arXiv:1606.06361v1 [cs.CL] 20 Jun 2016

A Probabilistic Generative Grammar for Semantic Parsing

Abulhair Saparov*
Carnegie Mellon University

Tom M. Mitchell*
Carnegie Mellon University

We present a framework that couples the syntax and semantics of natural language sentences in a generative model, in order to develop a semantic parser that jointly infers the syntactic, morphological, and semantic representations of a given sentence under the guidance of background knowledge. To generate a sentence in our framework, a semantic statement is first sampled from a prior, such as from a set of beliefs in a knowledge base. Given this semantic statement, a grammar probabilistically generates the output sentence. A joint semantic-syntactic parser is derived that returns the k-best semantic and syntactic parses for a given sentence. The semantic prior is flexible, and can be used to incorporate background knowledge during parsing, in ways unlike previous semantic parsing approaches. For example, semantic statements corresponding to beliefs in a knowledge base can be given higher prior probability, type-correct statements can be given somewhat lower probability, and beliefs outside the knowledge base can be given lower probability. The construction of our grammar invokes a novel application of hierarchical Dirichlet processes (HDPs), which in turn, requires a novel and efficient inference approach. We present experimental results showing, for a simple grammar, that our parser outperforms a state-of-the-art CCG semantic parser and scales to knowledge bases with millions of beliefs.

1. Introduction

Accurate and efficient semantic parsing is a long-standing goal in natural language processing. There are countless applications for methods that provide deep semantic analyses of sentences. Leveraging semantic information in text may provide improved algorithms for many problems in NLP, such as named entity recognition (Finkel and Manning 2009, 2010; Kazama and Torisawa 2007), word sense disambiguation (Tanaka et al. 2007; Bordes et al. 2012), semantic role labeling (Merlo and Musillo 2008), coreference resolution (Ponzetto and Strube 2006; Ng 2007), etc. A sufficiently expressive semantic parser may directly provide the solutions to many of these problems. Lower-level language processing tasks, such as those mentioned, may even benefit by incorporating semantic information, especially if the task can be solved jointly during semantic parsing.

Knowledge plays a critical role in natural language understanding. The formalisms used by most semantic parsing approaches require an ontology of entities and predicates, with which the semantic content of sentences can be represented. Moreover, even seemingly trivial sentences may have a large number of ambiguous interpretations. Consider the sentence *"She started the machine with the GPU,"* for example. Without additional knowledge, such as the fact that "machine" can refer to computing devices

* Machine Learning Department, 5000 Forbes Avenue, Pittsburgh, PA 15213.
E-mails: asaparov@cs.cmu.edu and tom.mitchell@cmu.edu

© 2005 Association for Computational Linguistics

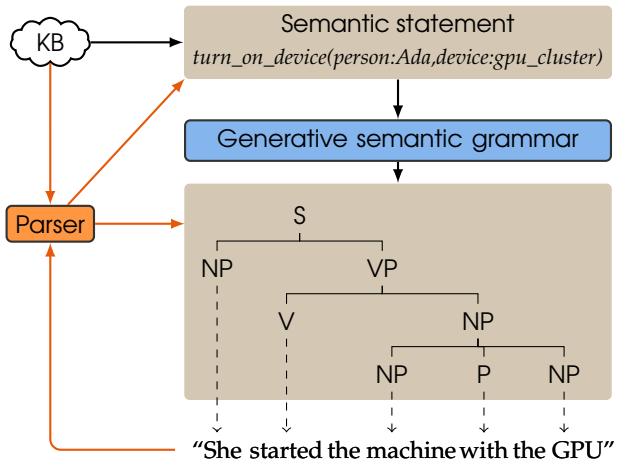

Figure 1
High-level illustration of the setting in which our grammar is applied. During parsing, the input is the observed sentence and knowledge base, and we want to find the k most probable semantic-syntactic parses given this input and the training data.

that contain GPUs, or that computers generally contain devices such as GPUs, the reader cannot determine whether the GPU is part of the machine or if the GPU is a device that is used to start machines.

The thesis underlying our research is that natural language understanding requires a belief system; that is, a large set of pre-existing beliefs related to the domain of discourse. Clearly, young children have many beliefs about the world when they learn language, and in fact, the process of learning language is largely one of learning to ground the meanings of words and sentences in these non-linguistically acquired beliefs. In some ways, the idea that language understanding requires a belief system is not new, as natural language researchers have been saying for years that background knowledge is essential to reducing ambiguity in sentence meanings (Bloom 2000; Anderson and Pearson 1984; Fincher-Kiefer 1992; Adams, Bell, and Perfetti 1995). But despite this general acknowledgement of the importance of background knowledge, we see very few natural language understanding systems that actually employ a large belief system as the basis for comprehending sentence meanings, and for determining whether the meaning of a new sentence contradicts, extends, or is already present in its belief system.

We present here a step in this direction: a probabilistic semantic parser that uses a large knowledge base (NELL) to form a prior probability distribution on the meanings of sentences it parses, and that "understands" each sentence either by identifying its existing beliefs that correspond to the sentence's meaning, or by creating new beliefs. More precisely, our semantic parser corresponds to a probabilistic generative model that assigns high probability to sentence semantic parses resulting in beliefs it already holds, lower prior probability to parses resulting in beliefs it does not hold but which are consistent with its more abstract knowledge about semantic types of arguments to different relations, and still lower prior probability to parses that contradict its beliefs about which entity types can participate in which relations.

This work is only a first step. It is limited in that we currently use it to parse sentences with a simple noun-verb-noun syntax (e.g. "Horses eat hay."), and considers only factual assertions in declarative sentences. Its importance is that it introduces a novel approach in which the semantic parser (a) prefers sentence semantic parses that

yield assertions it already believes, while (b) still allowing with lower prior probability sentence interpretations that yield new beliefs involving novel words, and (c) even allowing beliefs inconsistent with its background knowledge about semantic typing of different relations. We introduce algorithms for training the probabilistic grammar and producing parses with high posterior probability, given its prior beliefs and a new sentence. We present experimental evidence of the success and tractability of this approach for sentences with simple syntax, and evidence showing that the incorporated belief system, containing millions of beliefs, allows it to outperform state-of-the-art semantic parsers that do not hold such beliefs. Thus, we provide a principled, probabilistic approach to using a current belief system to guide semantic interpretation of new sentences which, in turn, can be used to augment and extend the belief system. We also argue that our approach can be extended to use the document-level context of a sentence as an additional source of background beliefs.

For reasons including but not limited to performance and complexity, most modern parsers operate over tokens, such as words. While this has worked sufficiently well for many applications, this approach assumes that a tokenization preprocessing step produces the correct output. This is nontrivial in many languages, such as Chinese, Thai, Japanese, and Tibetic languages. In addition, a large portion of the English vocabulary is created from the combination of simpler morphemes, such as the words "build-er," "in-describ-able," "anti-modern-ist." Moreover, language can be very noisy. Text messages, communication in social media, and real-world speech are but a few examples of noise obfuscating language. Standard algorithms for tokenization, lemmatization, and other preprocessing are oblivious to the underlying semantics, much less any background knowledge. Incorporating these components into a "joint parsing" framework will enable semantics and background knowledge to *jointly inform* lower-level processing of language. Our method couples semantics with syntax and other lower-level aspects of language, and can be guided by background knowledge via the semantic prior. We will demonstrate how this can be leveraged in our framework to model the morphology of individual verbs in a temporally-scoped relation extraction task.

Semantic statements are the logical expressions that represent meaning in sentences. For example, the semantic statement *turn_on_device(person:Ada, device:gpu_cluster)* may be used to express the meaning of the sentence example given earlier. There are many languages or *semantic formalisms* that can be used to encode these logical forms: first-order logic with lambda calculus (Church 1932), frame semantics (Baker, Fillmore, and Lowe 1998), abstract meaning representation (Banarescu et al. 2013), dependency-based compositional semantics (Liang, Jordan, and Klein 2013), vector-space semantics (Salton 1971; Turney and Pantel 2010), for example. Our approach is flexible and does not require the use of a specific semantic formalism.

In section 3, we review HDPs and describe the setting that we require to define our grammar. We present our approach in section 3.1.1 to perform HDP inference in this new setting. In section 4, we present the main generative process in our framework, and detail our application of the HDP. Although we present our model from a generative perspective, we show in the description of the framework that discriminative techniques can be integrated. Inference in our model is described in section 5. There, we present a chart-driven agenda parser that can leverage the semantic prior to guide its search. Finally, in section 6, we evaluate our parser on two relation-extraction tasks: the first is a task to extract simple predicate-argument representations from SVO sentences, and the second is a temporally-scoped relation extraction task that demonstrates our parser's ability to model the morphology of individual words, leading to improved generalization performance over words. Moreover, we demonstrate that the inclusion

of background knowledge from a knowledge base improves parsing performance on these tasks. The key contributions of this article are:

1. a framework to define grammars with coupled semantics, syntax, morphology, etc.,
2. the use of a prior on the semantic statement to incorporate prior knowledge,
3. and an efficient and exact k-best parsing algorithm guided by a belief system.

2. Background

Our model is an extension of context-free grammars (CFGs) (Chomsky 1956) that couples syntax and semantics. To generate a sentence in our framework, the semantic statement is first drawn from a prior. A grammar then recursively constructs a syntax tree top-down, randomly selecting production rules from distributions that depend on the semantic statement. We present a particular incarnation of a grammar in this framework, where hierarchical Dirichlet processes (HDPs) (Teh et al. 2006) are used to select production rules randomly. The application of HDPs in our setting is novel, requiring a new inference technique.

The use of the term "generative" does not refer to the Chomskian tradition of generative grammar (Chomsky 1957), although our approach does fall broadly within that framework. Rather, it refers to the fact that our model posits a probabilistic mechanism by which sentences are generated (by the *speaker*). Performing probabilistic inference under this model yields a parsing algorithm (the *listener*). This generative approach to modeling grammar underscores the duality between language generation and language understanding.

Our grammar can be related to synchronous CFGs (SCFGs) (Aho and Ullman 1972), which have been extended to perform semantic parsing (Li et al. 2015; Wong and Mooney 2007, 2006). However, in established use, SCFGs describe the generation of the syntactic and semantic components of sentences simultaneously, which makes the assumption that the induced probability distributions of the semantic and syntactic components factorize in a "parallel" manner. Our model instead describes the generation of the semantic component as a step with occurs *prior* to the syntactic component. This can be captured in SCFGs as a prior on the semantic start symbol, making no factorization assumptions on this prior. This is particularly useful when employing richer prior distributions on the semantics, such as a model of context or a knowledge base.

Adaptor grammars (Johnson, Griffiths, and Goldwater 2007) provide a framework that can jointly model the syntactic structure of sentences in addition to the morphologies of individual words (Johnson and Demuth 2010). Unlike previous work with adaptor grammars, our method couples syntax with semantics, and can be guided by background knowledge via the semantic prior. We will demonstrate how this can be leveraged in our framework to model the morphology of individual verbs in a temporally-scoped relation extraction task. Cohen, Blei, and Smith (2010) show how to perform dependency grammar induction using adaptor grammars. While grammar induction in our framework constitutes an interesting research problem, we do not address it in this work.

As in other parsing approaches, an equivalence can be drawn between our parsing problem and the problem of finding shortest paths in hypergraphs (Klein and Manning 2001, 2003a; Pauls and Klein 2009; Pauls, Klein, and Quirk 2010; Gallo, Longo, and Pallottino 1993). Our algorithm can then be understood as an application of A* search for the k-best paths in a very large hypergraph.

Our parser incorporates prior knowledge to guide its search, such as from an ontology and the set of beliefs in a knowledge base. Using this kind of approach, the parser can be biased to find context-appropriate interpretations in otherwise ambiguous or terse utterances. While systems such as Durrett and Klein (2014), Nakashole and Mitchell (2015), Kim and Moldovan (1995), and Salloum (2009) use background knowledge about the semantic types of different noun phrases to improve their ability to perform entity linking, co-reference resolution, prepositional phrase attachment, information extraction, and question answering, and systems such as Ratinov and Roth (2012), Durrett and Klein (2014), and Prokofyev et al. (2015) link noun phrases to Wikipedia entries to improve their ability to resolve co-references, these uses of background knowledge remain fragmentary. Krishnamurthy and Mitchell (2014) developed a CCG parser that incorporates background knowledge from a knowledge base during training through distant supervision, but their method is not able to do so during parsing. Our parser can be trained once, and then applied to a variety of settings, each with a different context or semantic prior.

3. Hierarchical Dirichlet processes

A core component of our statistical model is the *Dirichlet process* (DP) (Ferguson 1973), which can be understood as a distribution over probability distributions. If a distribution G is drawn from a DP, we can write $G \sim \text{DP}(\alpha, H)$, where the DP is characterized by two parameters: a concentration parameter $\alpha > 0$ and a base distribution H. The DP has the useful property that $\mathbb{E}[G] = H$, and the concentration parameter α describes the "closeness" of G to the base distribution H. In typical use, a number of parameters θ_i are drawn from a discrete distribution G, which is itself drawn from a Dirichlet process. The observations y_i are drawn using the parameters θ_i from another distribution F. This may be written as:

$$G \sim \text{DP}(\alpha, H), \tag{1}$$

$$\theta_1, \ldots, \theta_n \sim G, \tag{2}$$

$$y_i \sim F(\theta_i), \tag{3}$$

for $i = 1, \ldots, n$. In our application, we will define H to be a finite Dirichlet distribution and F is a categorical distribution. G can be marginalized out in the model above, resulting in the *Chinese restaurant process* representation (Aldous 1985):

$$\phi_1, \phi_2, \ldots \sim H, \tag{4}$$

$$z_i = \begin{cases} j & \text{with probability } \frac{\#\{k<i : z_k = j\}}{\alpha + i - 1}, \\ j^{new} & \text{with probability } \frac{\alpha}{\alpha + i - 1}, \end{cases} \tag{5}$$

$$\theta_i = \phi_{z_i} \text{ for } i = 1, \ldots, n, \tag{6}$$

$$y_i \sim F(\theta_i), \tag{7}$$

where $z_1 = 1$, $j^{new} = \max\{z_1, \ldots, z_{i-1}\} + 1$ is the indicator of a new table, and the quantity $\#\{k < i : z_k = j\}$ is the number of observations that were assigned to table j. The analogy is to imagine a restaurant where customers enter one at a time. Each customer chooses to sit at table j with probability proportional to the number of people

currently sitting at table j, or at a new table j^{new} with probability proportional to α. The i^{th} customer's choice is represented as z_i. As shown in later sections, this representation of the DP is amenable to inference using *Markov chain Monte Carlo* (MCMC) methods (Gelfand and Smith 1990; Robert and Casella 2010).

The *hierarchical Dirichlet process* (HDP) is an extension of the Dirichlet process for use in hierarchical modeling (Teh et al. 2006). An advantage of this approach is that statistical strength can be shared across nodes that belong to the same subtree. In an HDP, every node $\mathbf{n}$ in a fixed tree T is associated with a distribution $G^{\mathbf{n}}$, and:

$$G^{\mathbf{0}} \sim \text{DP}(\alpha^{\mathbf{0}}, H), \tag{8}$$

$$G^{\mathbf{n}} \sim \text{DP}(\alpha^{\mathbf{n}}, G^{\pi(\mathbf{n})}), \tag{9}$$

where $\pi(\mathbf{n})$ is the parent node of $\mathbf{n}$, and $\mathbf{0}$ is the root of T. In our application, the base distribution at the root H is Dirichlet. We can draw observations $y_1, \ldots, y_n$ from the HDP, given a sequence $x_1, \ldots, x_n$ of n paths from the root $\mathbf{0}$ to a leaf:

$$\theta_i \sim G^{x_i}, \tag{10}$$

$$y_i \sim F(\theta_i), \tag{11}$$

for $i = 1, \ldots, n$. For notational brevity, we write this equivalently as $y_i \sim \text{HDP}(x_i, T)$.

Just as marginalizing the Dirichlet process yields the Chinese restaurant process, marginalizing the HDP yields the *Chinese restaurant franchise* (CRF). For every node in the HDP tree $\mathbf{n} \in T$, there is a "Chinese restaurant" consisting of an infinite number of tables. Every table i in this restaurant at node $\mathbf{n}$ is assigned to a table in the *parent restaurant*. The assignment variable $z_i^{\mathbf{n}}$ is the index of the parent table to which table i in node $\mathbf{n}$ is assigned.

$$\phi_1^{\mathbf{0}}, \phi_2^{\mathbf{0}}, \ldots \sim H, \tag{12}$$

$$\text{for every node } \mathbf{n} \in T, \qquad z_i^{\mathbf{n}} = \begin{cases} j & \text{with probability } \propto n_j^{\pi(\mathbf{n})}, \\ j^{new} & \text{with probability } \propto \alpha^{\pi(\mathbf{n})}, \end{cases} \tag{13}$$

$$\phi_i^{\mathbf{n}} = \phi_{z_i^{\mathbf{n}}}^{\pi(\mathbf{n})}, \tag{14}$$

where $\pi(\mathbf{n})$ is the parent of node $\mathbf{n}$, and $n_j^{\pi(\mathbf{n})}$ is the current number of customers at node $\pi(\mathbf{n})$ sitting at table j. We are mildly abusing notation here, since $n_j^{\pi(\mathbf{n})}$ and $n^{\pi(\mathbf{n})}$ refer to the number of customers *at the time* $z_i^{\mathbf{n}}$ *is drawn* (which increases as additional $z_i^{\mathbf{n}}$ are drawn). To draw the observation y_i, we start with the leaf node at the end of the path x_i:

$$\theta_i = \phi_k^{x_i}, \tag{15}$$

$$y_i \sim F(\theta_i), \tag{16}$$

where $k - 1 = \#\{j < i : x_j = x_i\}$ is the number of previous observations drawn from node x_i.

3.1 Inference

In this section, we describe our method for performing posterior inference in the HDP. Let $\boldsymbol{z} = \{z_i^{\mathbf{n}} : \mathbf{n} \in T, i = 1, 2, \ldots\}$ be the set of table assignment variables in the HDP. If the distributions H and F are conditionally conjugate, as they are in our application, the $\boldsymbol{\phi}$ variables can be integrated out in closed form:

$$p(\boldsymbol{z}|\boldsymbol{x}, \boldsymbol{y}) = p(\boldsymbol{x})p(\boldsymbol{z}) \int p(\boldsymbol{y}|\boldsymbol{x}, \boldsymbol{z}, \boldsymbol{\phi}) d\boldsymbol{\phi}. \tag{17}$$

The posterior $p(\boldsymbol{z}|\boldsymbol{x}, \boldsymbol{y})$ is intractable to compute exactly, and so we approximate it by sampling. We obtain samples from $\boldsymbol{z}|\boldsymbol{x}, \boldsymbol{y}$ by performing collapsed Gibbs sampling as described in section 5.1 of Teh et al. (2006): we repeatedly sample $\boldsymbol{z}$ from its conditional distribution, with $\boldsymbol{\phi}$ integrated out:

$$z_i^{\mathbf{n}}|\boldsymbol{x}, \boldsymbol{y}, z_{-i}^{\mathbf{n}} = \begin{cases} j & \text{with prob.} \propto \#\{k \neq i : z_k^{\mathbf{n}} = j\} \cdot p(y_i^{\mathbf{n}}|\boldsymbol{x}, y_{-i}^{\mathbf{n}}, z_{-i}^{\mathbf{n}}, z_i^{\mathbf{n}} = j), \\ j^{new} & \text{with prob.} \propto \alpha^{\mathbf{n}} \cdot p(y_i^{\mathbf{n}}|\boldsymbol{x}, y_{-i}^{\mathbf{n}}, z_{-i}^{\mathbf{n}}, z_i^{\mathbf{n}} = j^{new}), \end{cases} \tag{18}$$

where $y_i^{\mathbf{n}}$ is the set of "descendant" observations of table i in node $\mathbf{n}$ (this includes observations assigned directly to the table, in addition to those assigned to tables further down in the hierarchy which themselves are assigned to this table), $y_{-i}^{\mathbf{n}} = \boldsymbol{y} \setminus y_i^{\mathbf{n}}$ is the set of all other observations, and $z_{-i}^{\mathbf{n}} = \boldsymbol{z} \setminus z_i^{\mathbf{n}}$ is the set of all other table assignment variables. Computing $p(y_i^{\mathbf{n}}|\boldsymbol{x}, y_{-i}^{\mathbf{n}}, z_{-i}^{\mathbf{n}}, z_i^{\mathbf{n}} = j)$ is straightforward since we can follow the chain of table assignments to the root. Let $r_i^{\mathbf{n}}$ be the *root cluster assignment* of the table i at node $\mathbf{n}$. In fact, we found it advantageous for performance to keep track of the root cluster assignments $\boldsymbol{r}$ for every table in the hierarchy. Thus, when $z_i^{\mathbf{n}} = j$, it must be the case that $y_i^{\mathbf{n}}$ were drawn from F with parameter $\phi^{\mathbf{0}}_{r_j^{\pi(\mathbf{n})}}$.

Computing $p(y_i^{\mathbf{n}}|\boldsymbol{x}, y_{-i}^{\mathbf{n}}, z_{-i}^{\mathbf{n}}, z_i^{\mathbf{n}} = j^{new})$ requires marginalizing over the assignment of the new table $z_{j^{new}}^{\pi(\mathbf{n})}$:

$$\begin{aligned} p(y_i^{\mathbf{n}}|\boldsymbol{x}, y_{-i}^{\mathbf{n}}, z_{-i}^{\mathbf{n}}, z_i^{\mathbf{n}} = j^{new}) = &\sum_{k=1}^{m^{\pi(\mathbf{n})}} \frac{n_k^{\pi(\mathbf{n})}}{n^{\pi(\mathbf{n})} + \alpha^{\pi(\mathbf{n})}} p(y_i^{\mathbf{n}}|\boldsymbol{x}, y_{-i}^{\mathbf{n}}, z_{-i}^{\mathbf{n}}, z_{j^{new}}^{\pi(\mathbf{n})} = k) \\ &+ \frac{\alpha^{\pi(\mathbf{n})}}{n^{\pi(\mathbf{n})} + \alpha^{\pi(\mathbf{n})}} p(y_i^{\mathbf{n}}|\boldsymbol{x}, y_{-i}^{\mathbf{n}}, z_{-i}^{\mathbf{n}}, z_{j^{new}}^{\pi(\mathbf{n})} = k^{new}), \end{aligned} \tag{19}$$

where $m^{\pi(\mathbf{n})}$ is the number of occupied tables at the node $\pi(\mathbf{n})$. At the root node $\pi(\mathbf{n}) = \mathbf{0}$, the above probability is just the prior of $y_i^{\mathbf{n}}$. We observe that the above probabilities are linear functions of the likelihoods $p(y_i^{\mathbf{n}}|\boldsymbol{x}, y_{-i}^{\mathbf{n}}, z_{-i}^{\mathbf{n}}, r_i^{\mathbf{n}} = k)$ for various root cluster assignments $r_i^{\mathbf{n}} = k$. Implemented naively, generating a single sample from equation 18 can take time linear in the number of clusters at the root, which would result in a quadratic-time algorithm for a single Gibbs iteration over all $\boldsymbol{z}$. However, we can exploit sparsity in the root cluster assignment likelihoods to improve performance. When $H = \text{Dir}(\beta)$ is a Dirichlet distribution and F is a categorical, then the collapsed root cluster

assignment likelihood is:

$$p(y_i^{\mathbf{n}}|\boldsymbol{x}, y_{-i}^{\mathbf{n}}, z_{-i}^{\mathbf{n}}, r_i^{\mathbf{n}} = k) = \frac{\prod_t \left(\beta_t + \#\{t \in y_k^{\mathbf{0}}\}\right)^{(\#\{t \in y_i^{\mathbf{n}}\})}}{\left(\sum_t \beta_t + \#y_k^{\mathbf{0}}\right)^{(\#y_i^{\mathbf{n}})}}. \tag{20}$$

Here, $a^{(b)}$ is the rising factorial $a(a+1)(a+2)\dots(a+b-1) = \frac{\Gamma(a+b)}{\Gamma(a)}$, and $\#\{t \in y_i^{\mathbf{n}}\}$ is the number of elements in $y_i^{\mathbf{n}}$ with value t. Notice that the denominator depends only on the sizes and not on the contents of $y_i^{\mathbf{n}}$ and $y_k^{\mathbf{0}}$. Caching the denominator values for common sizes of $y_i^{\mathbf{n}}$ and $y_k^{\mathbf{0}}$ can allow the sampler to avoid needless recomputation. This is especially useful in our application since many of the tables at the root tend to be small. Similarly, observe that the numerator factor is 1 for values of t where $\#\{t \in y_i^{\mathbf{n}}\} = 0$. Thus, the time required to compute the above probability is linear in the number of unique elements of $y_i^{\mathbf{n}}$, which can improve the scalability of our sampler. We perform the above computations in log space to avoid numerical overflow.

3.1.1 Computing probabilities of paths. In previous uses of the HDP, the paths x_i are assumed to be fixed. For instance, in document modeling, the paths correspond to documents or predefined categories of documents. In our application, however, the paths may be *random*. In fact, we will later show that our parser heavily relies on the posterior predictive distribution over paths, where the paths correspond to semantic parses. More precisely, given a collection of training observations $\boldsymbol{y} = \{y_1, \dots, y_n\}$ with their paths $\boldsymbol{x} = \{x_1, \dots, x_n\}$, we want to compute the probability of a new path x^{new} given a new observation y^{new}:

$$p(x^{new}|y^{new}, \boldsymbol{x}, \boldsymbol{y}) \propto p(x^{new}) \int p(y^{new}|\boldsymbol{z}, x^{new}) p(\boldsymbol{z}|\boldsymbol{x}, \boldsymbol{y}) d\boldsymbol{z}, \tag{21}$$

$$\approx \frac{p(x^{new})}{N_{samples}} \sum_{\boldsymbol{z}^* \sim \boldsymbol{z}|\boldsymbol{x},\boldsymbol{y}} p(y^{new}|\boldsymbol{z}^*, x^{new}). \tag{22}$$

Once we have the posterior samples $\boldsymbol{z}^*$, we can compute the quantity $p(y^{new}|\boldsymbol{z}^*, x^{new})$ by marginalizing over the table assignment for the new observation y:

$$p(y^{new}|\boldsymbol{z}^*, x^{new}) = \sum_{j=1}^{m^{x^{new}}} \frac{n_j^{x^{new}}}{n^{x^{new}} + \alpha^{x^{new}}} \, p(y^{new}|\boldsymbol{z}^*, \theta^{new} = \phi_j^{x^{new}}) + \frac{\alpha^{x^{new}}}{n^{x^{new}} + \alpha^{x^{new}}} \, p(y^{new}|\boldsymbol{z}^*, \theta^{new} = \phi_{j^{new}}^{x^{new}}). \tag{23}$$

Here, $m^{x^{new}}$ is the number of occupied tables at node x^{new}, $n_j^{x^{new}}$ is the number of customers sitting at table j at node x^{new}, and $n^{x^{new}}$ is the total number of customers at node x^{new}. The first term $p(y^{new}|\boldsymbol{z}^*, \theta^{new} = \phi_j^{x^{new}})$ can be computed since the j^{th} table exists and is assigned to a table in its parent node, which in turn is assigned to a table in its parent node, and so on. We can follow the chain of table assignments to the root. In the second term, the observation is assigned to a new table, whose assignment is unknown, and so we marginalize again over the assignment in the parent node for this

new table:

$$p(y^{new}|\boldsymbol{z}^*, \theta^{new} = \phi_{j^{new}}^{x^{new}}) = \sum_{j=1}^{m^{\pi(x^{new})}} \frac{n_j^{\pi(x^{new})}}{n^{\pi(x^{new})} + \alpha^{\pi(x^{new})}} \; p\left(y^{new} \middle| \boldsymbol{z}^*, \theta^{new} = \phi_j^{\pi(x^{new})}\right) + \frac{\alpha^{\pi(x^{new})}}{n^{\pi(x^{new})} + \alpha^{\pi(x^{new})}} \; p\left(y^{new} \middle| \boldsymbol{z}^*, \theta^{new} = \phi_{j^{new}}^{\pi(x^{new})}\right), \quad (24)$$

where $\pi(x^{new})$ is the parent node of x^{new}. Again, the probability in the first term can be computed as before, but the probability in the second term depends on the assignment of the new table, which is unknown. Thus, since it is possible that a new table will be created at every level in the hierarchy up to the root, we can apply this formula recursively. At the root $\mathbf{0}$, the probability $p(y^{new}|\boldsymbol{z}^*, \theta^{new} = \phi_{j^{new}}^{\mathbf{0}})$ is just the prior probability of y^{new}.

If the tree T is small, it is straightforward to compute the quantity in equation 22 for every path x^{new} in the tree, using the method described above. In our application however, the size of T depends on the size of the ontology, and may easily become very large. In this case, the naïve approach becomes computationally infeasible. As such, we develop an algorithm to incrementally find the k best paths that maximize the quantity in equation 22. For sparse distributions, where most of the probability mass is concentrated in a small number of paths x^{new}, this algorithm can effectively characterize the predictive distribution in equation 21. The algorithm is essentially a search over nodes in the tree, starting at the root and descending the nodes of the tree T, guided through paths of high probability. Each search state $\mathtt{s}$ consists of the following fields:

- $\mathtt{s.n}$ is the current position of the search in the tree.
- $\mathtt{s.v}$ is an array of probability scores of length $N_{samples}$. Each element in this array represents the probability of drawing the observation y^{new} from the current node $\mathtt{s.n}$, and thus is identical to the probability of assigning y^{new} to a new table at any child node of $\mathtt{s.n}$. This is useful to compute the quantity in equation 22 using the recursive method as described above.

The search is outlined in algorithm 1. We observe that the quantity in equation 22 is a sum of independent functions, each being a linear combination of the terms $p(y^{new}|\boldsymbol{z}_i^*, \theta^{new} = \phi_j^{\mathbf{n}})$ over the tables available at node $\mathbf{n}$ and the new table $p(y^{new}|\boldsymbol{z}_i^*, \theta^{new} = \phi_{j^{new}}^{\mathbf{n}})$ (this latter probability is stored in $\mathtt{s.v}_i$). Thus, the upper bound on equation 22 over all paths that pass through node $\mathtt{s.n}$ is:

$$\max_{\{x^{new}:\mathtt{s.n}\in x^{new}\}} \frac{p(x^{new})}{N_{samples}} \sum_{i=1}^{N_{samples}} \max_{j=1,\ldots,m^{\mathtt{s.n}}} \left\{p(y^{new}|\boldsymbol{z}_i^*, \theta^{new} = \phi_j^{\mathtt{s.n}}), \mathtt{s.v}_i\right\}. \quad (25)$$

We sort elements in the priority queue using this expression.

As a result, once the algorithm has completed k items, we are guaranteed that the search has found k best paths. Thus, an "iterator" data structure can be efficiently implemented using this algorithm, which returns paths x^{new} in order of decreasing predictive probability, with the first item being optimal. The search algorithm can be modified for other representations of the HDP, and can be extended to the case where H and F are not conjugate. It may also be incorporated into a larger inference procedure to jointly infer the paths $\boldsymbol{x}$ and the latent variables in the HDP. It is also straightforward to compute predictive probabilities where the path x^{new} is restricted to a subset of paths X:

Algorithm 1: Search algorithm to find the k best paths in the HDP that maximize the quantity in equation 22.

initialize priority queue with initial state $\texttt{s}$
$\texttt{s.n} \leftarrow \mathbf{0}$ `/* start at the root */`
for $i = 1, \ldots, N_{samples}$, **do**
  $\texttt{s.v}_i \leftarrow \sum_{j=1}^{m^{\mathbf{0}}} \frac{n_j^{\mathbf{0}}}{n^{\mathbf{0}}+\alpha^{\mathbf{0}}} p(y^{new}|\boldsymbol{z}_i^*, \theta^{new} = \phi_j^{\mathbf{0}}) + \frac{\alpha^{\mathbf{0}}}{n^{\mathbf{0}}+\alpha^{\mathbf{0}}} p(y^{new}|\boldsymbol{z}_i^*, \theta^{new} = \phi_{j^{new}}^{\mathbf{0}})$
repeat
  pop state $\texttt{s}$ from the priority queue
  if $\texttt{s.n}$ *is a leaf*
    complete the path $\texttt{s.n}$ with probability $\frac{p\{x^{new}=\texttt{s.n}\}}{N_{samples}} \sum_{i=1}^{N_{samples}} \texttt{s.v}_i$
  foreach *child node* $\mathbf{c}$ *of* $\texttt{s.n}$, **do**
    create new search state $\texttt{s}^*$
    $\texttt{s}^*\texttt{.n} \leftarrow \mathbf{c}$
    for $i = 1, \ldots, N_{samples}$, **do**
      $\texttt{s}^*\texttt{.v}_i \leftarrow \sum_{j=1}^{m^{\mathbf{c}}} \frac{n_j^{\mathbf{c}}}{n^{\mathbf{c}}+\alpha^{\mathbf{c}}} p(y^{new}|\boldsymbol{z}_i^*, \theta^{new} = \phi_j^{\mathbf{c}}) + \frac{\alpha^{\mathbf{c}}}{n^{\mathbf{c}}+\alpha^{\mathbf{c}}} \texttt{s.v}_i$
    push $\texttt{s}^*$ onto priority queue with key in equation 25
until *there are* k *completed paths*

$p(x^{new}|y^{new}, \boldsymbol{x}, \boldsymbol{y}, x^{new} \in X)$. To do so, the algorithm is restricted to only expand nodes that belong to paths in X.

An important concern when performing inference with very large trees T is that it is not feasible to explicitly store every node in memory. Fortunately, collapsed Gibbs sampling does not require storing nodes whose descendants have zero observations. In addition, algorithm 1 can be augmented to avoid storing these nodes, as well. To do so, we make the observation that for any node $\mathbf{n} \in T$ in the tree whose descendants have no observations, $\mathbf{n}$ will have zero occupied tables. Therefore, the probability $p(y^{new}|\boldsymbol{z}^*, x^{new}) = p(y^{new}|\boldsymbol{z}^*, \theta^{new} = \phi_{j^{new}}^{\mathbf{n}})$ is identical for any path x^{new} that passes through $\mathbf{n}$. Thus, when the search reaches node $\mathbf{n}$, it can simultaneously complete all paths x^{new} that pass through $\mathbf{n}$, and avoid expanding nodes with zero observations among its descendants. As a result, we only need to explicitly store a number of nodes linear in the size of the training data, which enables practical inference with very large hierarchies.

There is a caveat that arises when we wish to compute a *joint predictive probability* $p(x_1^{new}, \ldots, x_k^{new}|y_1^{new}, \ldots, y_k^{new}, \boldsymbol{x}, \boldsymbol{y})$, where we have multiple novel observations. Re-writing equation 21 in this setting, we have:

$$p(x_1^{new}, \ldots, x_k^{new}|y_1^{new}, \ldots, y_k^{new}, \boldsymbol{x}, \boldsymbol{y}) \propto p(\boldsymbol{x}^{new}) \int p(y_1^{new}, \ldots, y_k^{new}|\boldsymbol{z}^*, \boldsymbol{x}^{new}) p(\boldsymbol{z}|\boldsymbol{x}, \boldsymbol{y}) d\boldsymbol{z}. \quad (26)$$

For the CRF, the joint likelihood $p(y_1^{new}, \ldots, y_k^{new}|\boldsymbol{z}^*, \boldsymbol{x}^{new})$ does not factorize, since the observations are not independent (they are exchangeable). One workaround is to use a representation of the HDP where the joint likelihood factorizes, such as the direct assignment representation (Teh et al. 2006). Another approach is to approximate the

joint likelihood with the factorized likelihood. In our parser, we instead make the following approximation:

$$p(y_1^{new}, \ldots, y_k^{new} | \boldsymbol{x}^{new}, \boldsymbol{x}, \boldsymbol{y}) = \prod_{i=1}^{k} p(y_i^{new} | y_1^{new}, \ldots, y_{i-1}^{new}, \boldsymbol{x}^{new}, \boldsymbol{x}, \boldsymbol{y}) \tag{27}$$

$$\approx \prod_{i=1}^{k} p(y_i^{new} | \boldsymbol{x}^{new}, \boldsymbol{x}, \boldsymbol{y}). \tag{28}$$

Substituting into equation 26, we obtain:

$$p(\boldsymbol{x}^{new} | \boldsymbol{y}^{new}, \boldsymbol{x}, \boldsymbol{y}) \propto p(\boldsymbol{x}^{new}) \prod_{i=1}^{k} \int p(y_i^{new} | \boldsymbol{z}^*, \boldsymbol{x}^{new}) p(\boldsymbol{z} | \boldsymbol{x}, \boldsymbol{y}) d\boldsymbol{z}. \tag{29}$$

When the size of the training data $(\boldsymbol{x}, \boldsymbol{y})$ is large with respect to the test data $(\boldsymbol{x}^{new}, \boldsymbol{y}^{new})$, the approximation works well, which we also find to be the case in our experiments.

4. Generative semantic grammar

We present a generative model of text sentences. In this model, semantic statements are generated probabilistically from some higher-order process. Given each semantic statement, a formal grammar selects text phrases, which are concatenated to form the output sentence. We present the model such that it remains flexible with regard to the semantic formalism. Even though our grammar can be viewed as an extension of context-free grammars, it is important to note that our model of grammar is only *conditionally context-free*, given the semantic statement. Otherwise, if the semantic information is marginalized out, the grammar is sensitive to context.

4.1 Definition

Let $\mathcal{N}$ be a set of nonterminals, and let $\mathcal{W}$ be a set of terminals. Let $\boldsymbol{R}$ be a set of production rules which can be written in the form $\mathrm{A} \to \mathrm{B}_1 \ldots \mathrm{B}_k$ where $\mathrm{A} \in \mathcal{N}$ and $\mathrm{B}_1, \ldots, \mathrm{B}_k \in \mathcal{W} \cup \mathcal{N}$. The tuple $(\mathcal{W}, \mathcal{N}, \boldsymbol{R})$ is a *context-free grammar* (CFG) (Chomsky 1956).

We couple syntax with semantics by augmenting the production rules $\boldsymbol{R}$. In every production rule $\mathrm{A} \to \mathrm{B}_1 \ldots \mathrm{B}_k$ in $\boldsymbol{R}$, we assign to every right-hand side symbol B_i a surjective operation $f_i : \mathcal{X}_A \mapsto \mathcal{X}_{B_i}$ that transforms semantic statements, where $\mathcal{X}_A$ is the set of semantic statements associated with the symbol A and $\mathcal{X}_{B_i}$ is the set of semantic statements associated with the symbol B_i. Intuitively, the operation describes how the semantic statement is "passed on" to the child nonterminals in the generative process. During parsing, these operations will describe how simpler semantic statements combine to form larger statements, enabling semantic compositionality. For example, suppose we have a semantic statement $x =$ *has_color(reptile:frog,color:green)* and the production rule $\mathrm{S} \to \mathrm{NP}\ \ \mathrm{VP}$. We can pair the semantic operation f_1 with the NP in the right-hand side such that $f_1(x) =$ *reptile:frog* selects the subject argument. Similarly, we can pair the semantic operation f_2 with the VP in the right-hand side such that $f_2(x) = x$ is the identity operation. The augmented production rule is $(\mathrm{A}, \mathrm{B}_1, \ldots, \mathrm{B}_k, f_1, \ldots, f_k)$

and the set of augmented rules is $\boldsymbol{R}^*$. In parsing, we require the computation of the inverse of semantic operations, which is the preimage of a given semantic statement $f^{-1}(x) = \{x' : f(x') = x\}$. Continuing the example above, f_1^{-1}(*reptile:frog*) returns a set that contains the statement *has_color(reptile:frog,color:green)* in addition to statements like *eats_insect(reptile:frog,insect:fly)*.

To complete the definition of our grammar, we need to specify the method that, given a nonterminal A $\in \mathcal{N}$ and a semantic statement $x \in \mathcal{X}_A$, selects a production rule from the set of rules in $\boldsymbol{R}^*$ with the left-hand side nonterminal A. To accomplish this, we define $\texttt{select}_{A,x}$ as a distribution over rules from $\boldsymbol{R}^*$ that has A as its left-hand side, dependent on x. We will later provide a number of example definitions of this $\texttt{select}_{A,x}$ distribution. Thus, a grammar in our framework is fully specified by the tuple $(\mathcal{W}, \mathcal{N}, \boldsymbol{R}^*, \texttt{select})$.

Note that other semantic grammar formalisms can be fit into this framework. For example, in categorical grammars, a lexicon describes the mapping from elementary components of language (such as words) to a syntactic category and a semantic meaning. Rules of inference are available to combine these lexical items into (tree-structured) derivations, eventually resulting in a syntactic and semantic interpretation of the full sentence (Steedman 1996; Jäger 2004). In our framework, we imagine this process in reverse. The set $\mathcal{X}_S$ is the set of all derivable semantic statements with syntactic category S. The generative process begins by selecting one statement from this set $x \in \mathcal{X}_S$. Next, we consider all applications of the rules of inference that would yield x, with each unique application of an inference rule being equivalent to a production rule in our framework. We select one of these production rules according to our generative process and continue recursively. The items in the lexicon are equivalent to preterminal production rules in our framework. Thus, the generative process below describes a way to endow parses in categorical grammar with a probability measure. This can be used, for example, to extend earlier work on generative models with CCG (Hockenmaier 2001; Hockenmaier and Steedman 2002). Different choices of the `select` distribution induce different probability distributions over parses.

We do not see a straightforward way to fit linear or log-linear models over full parses into our framework, where a vector of features can be computed for each full parse (Berger, Pietra, and Pietra 1996; Ratnaparkhi 1998). This is due to our assumption that, given the semantic statement, the probability of a parse factorizes over the production rules used to construct that parse. However, the `select` distribution can be defined using linear and log-linear models, as we will describe in section 4.3.

4.2 Generative process

The process for generating sentences in this framework begins by drawing a semantic statement $x \in \mathcal{X}_S$ where S is the root nonterminal. Thus, there is a prior distribution $p(x)$ for all $x \in \mathcal{X}_S$. Next, the syntax is generated top-down starting at S. We draw a production rule with S as the left-hand side from $\texttt{select}_{S,x}$. The semantic transformation operations f_i are applied to x and the process is repeated for the right-hand side nonterminals. More concretely, we define the following operation `expand` which takes two arguments: a symbol $\mathrm{A} \in \mathcal{W} \cup \mathcal{N}$ and a semantic statement $x \in \mathcal{X}_A$.

```
1 function expand(x, A)
2     if A ∈ 𝒲
          /* simply return the word if A is a terminal */
3         return A
4     else
          /* select a production rule with form A → B_1, ..., B_k */
5         (A, B_1, ..., B_k, f_1, ..., f_k) ~ select_{A,x}
6         return yield(expand(f_1(x), B_1), ..., expand(f_k(x), B_k))
```

The `yield` operation concatenates strings into a single output string. Then, the output sentence y is generated simply by $y = \texttt{expand}(x, \mathrm{S})$. Depending on the application, we may require that the generative process capitalizes the first letter of the output sentence, and/or appends terminating punctuation to the end. A noise model may also be appended to the generative process. The above algorithm may be easily extended to also return the full syntax tree.

4.3 Selecting production rules

There are many possible choices for the `select` distribution. The most straightforward is to define a categorical distribution over the available production rules, and simply draw the selected rule from this distribution. The result would be a simple extension of probabilistic context-free grammars (PCFGs) that couples semantics with syntax. However, this would remove any dependence between the semantic statement and the production rule selection.

To illustrate the importance of this dependence, consider generating a sentence with the semantic statement *athlete_plays_sport(athlete:roger_federer,sport:tennis)* using the grammar in figure 2 (the process is graphically depicted in figure 3). We start with the root nonterminal S:

step 1 We can only select the first production rule, and so we apply the semantic operation `select_arg1` on the semantic statement to obtain *athlete:roger_federer* for the right-hand side nonterminal N. We apply the semantic operation `delete_arg1` to obtain *athlete_plays_sport(·,sport:tennis)* for VP.

step 2 Expanding N, we `select` a terminal symbol given the semantic statement *athlete:roger_federer*. Suppose "Andre Agassi" is returned.

step 3 Now, we expand the VP symbol. We draw from $\texttt{select}_{\mathrm{VP}}$ to choose one of the two available production rules. Suppose the rule VP → V N is selected. Thus, we apply the `identity` operation for the V nonterminal to obtain

S → N : `select_arg1` VP : `delete_arg1`
VP → V : `identity` N : `select_arg2`
VP → V : `identity`

N → "tennis"
N → "Andre Agassi"
N → "Chopin"
V → "swims"
V → "plays"

Figure 2
Example of a grammar in our framework. This grammar operates on semantic statements of the form *predicate(first argument, second argument)*. The semantic operation `select_arg1` returns the first argument of the semantic statement. Likewise, the operation `select_arg2` returns the second argument. The operation `delete_arg1` removes the first argument, and `identity` returns the semantic statement with no change.

athlete_plays_sport(·,sport:tennis). We similarly apply `select_arg2` for the N nonterminal to obtain *sport:tennis*.

step 4 We expand the V nonterminal, drawing from $\texttt{select}_{\text{V}}$ on the semantic statement *athlete_plays_sport(·,sport:tennis)*. Suppose "plays" is returned.

step 5 Finally, we expand the N nonterminal, drawing from $\texttt{select}_{\text{N}}$ with the statement *sport:tennis*. Suppose "tennis" is returned. We concatenate all returned strings to form the sentence "Andre Agassi plays tennis."

However, now consider generating another sentence with the same grammar for the statement *athlete_plays_sport(athlete:roger_federer,* ***sport:swimming****)*. In *step 3* of the above process, the `select` distribution would necessarily have to depend on the semantic statement. In English, the probability of observing a sentence of the form N V N ('Rachmaninoff makes music') versus N V ('Rachmaninoff composes') depends on the underlying semantic statement.

To capture this dependence, we use HDPs to define the `select` distribution. Every nonterminal A $\in \mathcal{N}$ is associated with an HDP, and in order to fully specify the grammar, we need to specify the structure of each HDP tree. Let T_A be the tree associated with the nonterminal A. The model is flexible with how the trees are defined, but we construct trees with the following method. First, select m discrete features $g_1, \ldots, g_m$ where each $g_i : \mathcal{X} \mapsto \mathbb{Z}$ and $\mathbb{Z}$ is the set of integers. These features operate on *semantic statements*. For example, suppose we restrict the space of semantic statements to be the set of single predicate instances (triples). The relations in an ontology can be assigned unique integer indices, and so we may define a semantic feature as a function which simply returns the index of the predicate given a semantic statement. We construct the HDP tree T_A starting with the root, we add a child node for every possible output of g_1. We repeat the process recursively, constructing a complete tree of depth $m + 1$.

As an example, we will construct a tree for the nonterminal VP for the example grammar in figure 2. Suppose in our ontology, we have the predicates *athlete_plays_sport* and *musician_plays_instrument*, labeled 0 and 1, respectively. The ontology also contains the concepts *athlete:roger_federer*, *sport:tennis*, and *sport:swimming*, also labeled 0, 1, and 2, respectively. We define the first feature g_1 to return the predicate index. The second feature g_2 returns the index of the concept in the second argument of the semantic statement. The tree is constructed starting with the root, we add a child node for each predicate in the ontology: *athlete_plays_sport* and *musician_plays_instrument*. Next, for each child node, we add a grandchild node for every concept in the ontology: *athlete:roger_federer*, *sport:tennis*, and *sport:swimming*. The resulting tree T_{VP} has depth

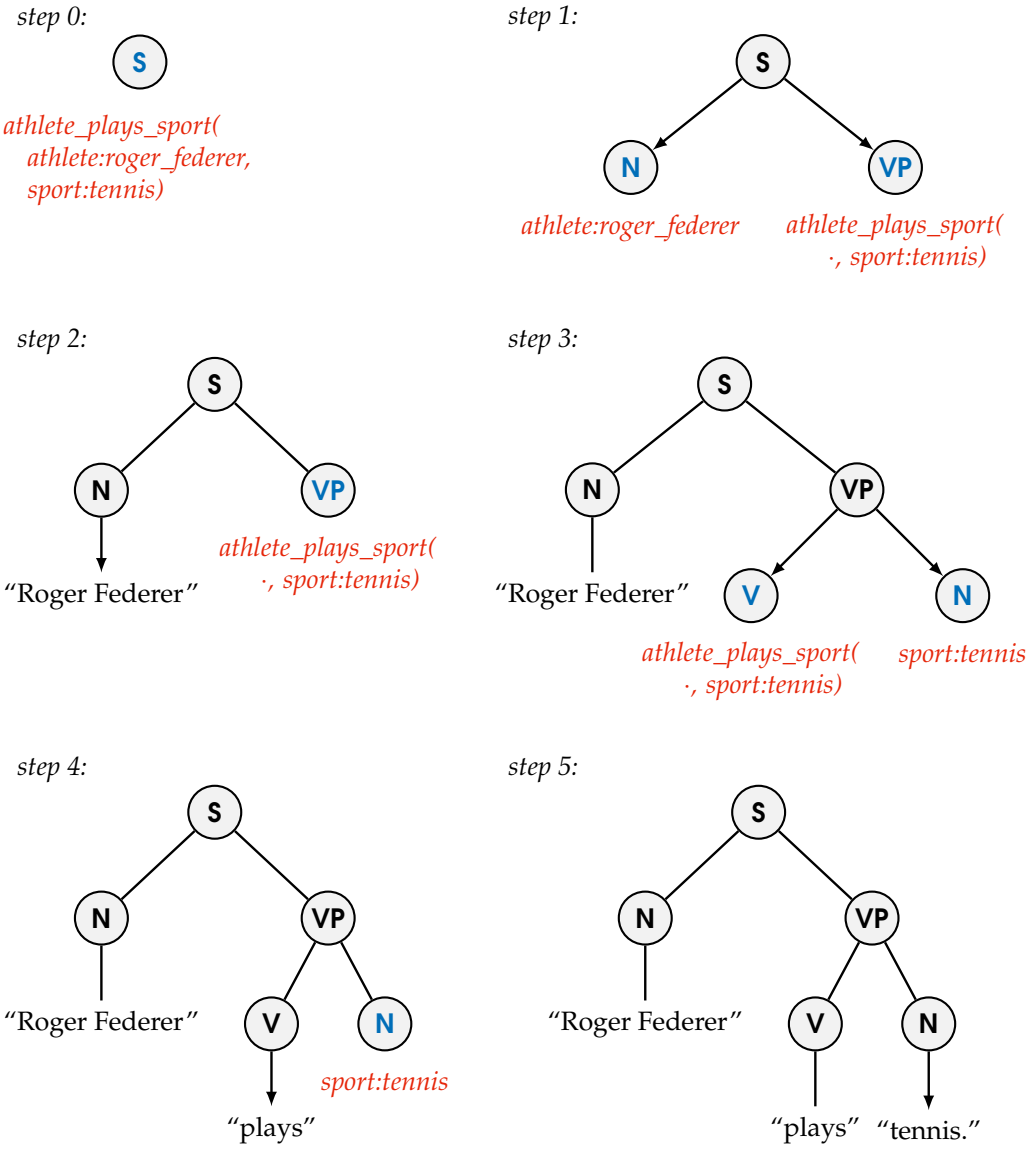

Figure 3
A depiction of the generative process producing a sentence for the semantic statement *athlete_plays_sport(athlete:roger_federer,sport:tennis)* using the grammar in figure 2.

2, with a root node with 2 child nodes, and each child node has 3 grandchild nodes. This construction enables the `select` distribution for the nonterminal VP to depend on the predicate and the second argument of the semantic statement.

With the fully-specified HDPs and their corresponding trees, we have fully specified `select`. When sampling from $\texttt{select}_{\text{A},x}$ for the nonterminal $\text{A} \in \mathcal{N}$ and a semantic statement $x \in \mathcal{X}$, we compute the m semantic features for the given semantic statement: $g_1(x), g_2(x), \ldots, g_m(x)$. This sequence of indices specifies a path from the root of the tree

down to a leaf. We then simply draw a production rule observation from this leaf node, and return the result: $r \sim \text{HDP}(x, T_A) = \texttt{select}_{\text{A},x}$.

There are many other alternatives for defining the `select` distribution. For instance, a log-linear model can be used to learn dependence on a set of features. The HDP provides statistical advantages, smoothing the learned distributions, resulting in a model more robust to data sparsity issues.

In order to describe inference in this framework, we must define additional concepts and notation. For a nonterminal $\text{A} \in \mathcal{N}$, observe that the paths from the root to the leaves of its HDP tree induce a partition on the set of semantic statements $\mathcal{X}_A$. More precisely, two semantic statements $x_1, x_2 \in \mathcal{X}_A$ belong to the same equivalence class if they correspond to the same path in an HDP tree.

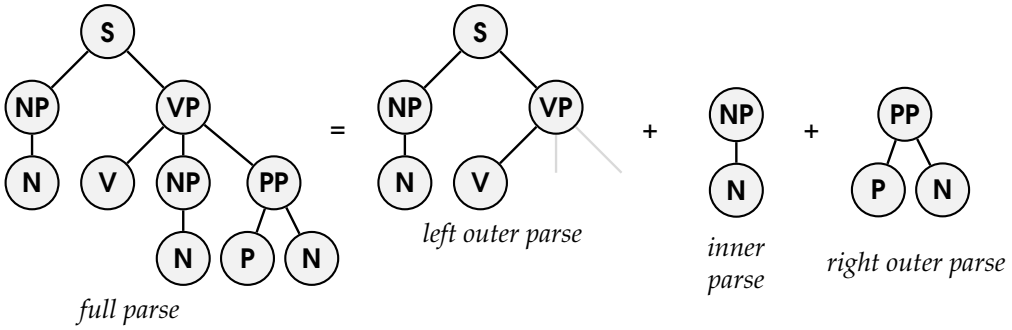

Figure 4
An example decomposition of a parse tree into its left outer parse, inner parse (of the object noun phrase), and its right outer parse. This is one example of such a decomposition. For instance, we may similarly produce a decomposition where the prepositional phrase is the inner parse, or where the verb is the inner parse. The terminals are omitted and only the syntactic portion of the parse is displayed here for consiseness.

Every parse (x, s) consists of a semantic statement x and a syntax tree s. The syntax tree s is a rooted tree containing an interior vertex for every nonterminal and a leaf for every terminal. Every vertex is associated with a start position and end position in the sentence. An interior vertex along with its immediate children corresponds to a particular production rule in the grammar $(\text{A} \to \text{B}_1{:}f_1 \dots \text{B}_n{:}f_n) \in \boldsymbol{R}^*$, where the interior vertex is associated with the nonterminal A and its children respectively correspond to the symbols $\text{B}_1, \dots, \text{B}_n$, left-to-right. Thus, every edge in the tree is labeled with a semantic transformation operation. A subgraph s_I of s can be called an *inner syntax tree*. The corresponding *outer syntax tree* s_O is $s_O = s \setminus s_I$ is the syntax tree with s_I deleted. We further draw a distinction between left and right components of an outer syntax tree. Define the *left outer syntax tree* s_L as the minimal subgraph of s_O containing all subtrees positioned to the left of s_I, and containing all ancestor vertices of s_I. The *right outer syntax tree* s_R forms the remainder of the outer parse, and so s can be decomposed into three distinct trees: $s = s_L \cup s_R \cup s_I$. See figure 4 for an illustration. Note that it is possible that s_R consists of multiple disconnected trees. In the description of our parser, we will frequently use the notation $p(s)$ to refer to the joint probability of all the production rules in the syntax tree s; that is, $p(s) = p(\bigcap_{(\text{A}\to\beta)\in s} \text{A} \to \beta)$, where β is the right-hand side of some production rule.

5. Inference

Let $\boldsymbol{y} \triangleq \{y_1, \ldots, y_n\}$ be a collection of training sentences, along with their corresponding syntax trees $\boldsymbol{s} \triangleq \{s_1, \ldots, s_n\}$ and semantic statement labels $\boldsymbol{x} \triangleq \{x_1, \ldots, x_n\}$. Given a new sentence y^{new}, the goal of parsing is to compute the probability of its semantic statement x^{new} and syntax s^{new}:

$$p(x^{new}, s^{new} | y^{new}, \boldsymbol{x}, \boldsymbol{s}, \boldsymbol{y}) \propto \int p(x^{new}, s^{new}, y^{new} | \boldsymbol{\theta}) p(\boldsymbol{\theta} | \boldsymbol{x}, \boldsymbol{s}, \boldsymbol{y}) d\boldsymbol{\theta}. \tag{30}$$

In this expression, $\boldsymbol{\theta}$ are the latent variables in the grammar. Different applications will rely on this probability in different ways. For example, we may be interested in the semantic parse that maximizes this probability. The above integral is intractable to compute exactly, so we use *Markov chain Monte Carlo* (MCMC) to approximate it:

$$\approx \frac{1}{N_{samples}} \sum_{\boldsymbol{\theta}^* \sim \boldsymbol{\theta} | \boldsymbol{x}, \boldsymbol{s}, \boldsymbol{y}} p(x^{new}, s^{new}, y^{new} | \boldsymbol{\theta}^*), \tag{31}$$

where the sum is taken over samples from the posterior of the latent grammar variables $\boldsymbol{\theta}$ given the training data $\boldsymbol{x}$, $\boldsymbol{s}$, and $\boldsymbol{y}$.[1]

We make the assumption that the likelihood factorizes over the nonterminals. More precisely:

$$p(y^{new}, s^{new} | x^{new}, \boldsymbol{\theta}) = \prod_{A \in \mathcal{N}} p(\{\mathrm{A} \to \beta \in s^{new}\} | x^{new}, \theta_A), \tag{32}$$

where θ_A are the latent variables specific to the nonterminal A, and $\{\mathrm{A} \to \beta \in s^{new}\}$ is the set of production rules in s^{new} that have A as the left-hand side nonterminal. Thus, we may factorize the joint likelihood as:

$$p(x^{new}, s^{new}, y^{new} | \boldsymbol{\theta}) = p(x^{new}) \prod_{\mathrm{A} \in \mathcal{N}} p\left(\{\mathrm{A} \to \beta \in s^{new}\} | x^{new}, \theta_A\right), \tag{33}$$

where the first product is over the nonterminals $\mathrm{A} \in \mathcal{N}$ in the grammar. Note that the probability $p\left(\mathrm{A} \to \beta | x^{new}, \theta_A\right)$ is equivalent to the probability of drawing the rule $\mathrm{A} \to \beta$ from $\texttt{select}_{\mathrm{A}, x^{new}}$ for nonterminal A and semantic statement x^{new}. Plugging

1 We also attempted a variational approach to inference, approximating the integral as $\mathbb{E}_q[p(x^{new}, s^{new}, y^{new} | \boldsymbol{\theta})]$, where q was selected to minimize the KL divergence to the posterior $p(\boldsymbol{\theta} | \boldsymbol{x}, \boldsymbol{s}, \boldsymbol{y})$. We experimented with a number of variational families, but we found that they were not sufficiently expressive to accurately approximate the posterior for our purposes.

equation 33 into 30 and 31, we obtain:

$$p(x^{new}, s^{new}|y^{new}, \boldsymbol{x}, \boldsymbol{s}, \boldsymbol{y})$$

$$\propto p(x^{new}) \prod_{\mathrm{A}\in\mathcal{N}} \int p\left(\{\mathrm{A}\rightarrow\beta\in s^{new}\}|x^{new}, \theta_A\right) p(\theta_A|\boldsymbol{x}, \boldsymbol{s}, \boldsymbol{y}) d\theta_A, \quad (34)$$

$$\approx \frac{p(x^{new})}{N_{samples}^{|\mathcal{N}|}} \prod_{\mathrm{A}\in\mathcal{N}} \prod_{(\mathrm{A}\rightarrow\beta)\in s^{new}} \sum_{\theta_A^*\sim\theta_A|\boldsymbol{x},\boldsymbol{s},\boldsymbol{y}} p\left(\mathrm{A}\rightarrow\beta|x^{new}, \theta_A^*\right), \quad (35)$$

where the second product iterates over the production rules that constitute the syntax s^{new}. Note that we applied the approximation as described in equation 28. The semantic prior $p(x^{new})$ plays a critically important role in our framework. It is through this prior that we can add dependence on background knowledge during parsing. Although we present a setting in which training is supervised with both syntax trees and semantic labels, it is straightforward to apply our model in the setting where we have semantic labels but syntax information is missing. In such a setting, a Gibbs step can be added where the parser is run on the input sentence with the fixed semantic statement, returning a distribution over syntax trees for each sentence.

Now, we divide the problem of inference into two major components:

Inference over HDP paths: Given a set of semantic statements $X \subseteq \mathcal{X}$, incrementally find the k best semantic statements $x \in X$ that maximize the sum $\sum p(\mathrm{A}\rightarrow\beta|x, \theta_A)$ within equation 35. We observe that this quantity only depends on the HDP associated with nonterminal A. Note that this is exactly the setting as described in section 3.1.1, and so we can directly apply algorithm 1 to implement this component.

Parsing: Efficiently compute the k most likely semantic and syntactic parses $\{x^{new}, s^{new}\}$ that maximize $p(x^{new}, s^{new}|y^{new}, \boldsymbol{x}, \boldsymbol{s}, \boldsymbol{y})$ for a given sentence y^{new}. We describe this component in greater detail in the next section. This component utilizes the previous component.

5.1 Parsing

We develop a top-down parsing algorithm that computes the k-best semantic/syntactic parses (x^{new}, s^{new}) that maximize $p(x^{new}, s^{new}|y^{new}, \boldsymbol{x}, \boldsymbol{s}, \boldsymbol{y})$ for a given sentence y^{new}. We emphasize that this parser is largely independent of the choice of the distribution `select`. The algorithm searches over a space of items called *rule states*, where each rule state represents the parser's position within a specific production rule of the grammar. *Complete rule states* represent the parser's position after completing parsing of a rule in the grammar. The algorithm also works with *nonterminal structures* that represent a completed parse of a nonterminal within the grammar. The parser keeps a priority queue of unvisited rule states called the *agenda*. A data structure called the *chart* keeps intermediate results on contiguous portions of the sentence. A predefined set of operations are available to the algorithm. At every iteration of the main loop, the algorithm pops the rule state with the highest weight from the agenda and adds it to the chart, applying any available operation on this state using any intermediate structures in the chart. These operations may add additional rule states to the agenda, with priority given by an upper bound on $\log p(x^{new}, s^{new}|y^{new}, \boldsymbol{x}, \boldsymbol{s}, \boldsymbol{y})$. The overall structure of our parser is reminiscent of the Earley parsing algorithm, which is the classical example of

a top-down parsing algorithm for CFGs (Earley 1970). We will draw similarities in our description below. The parsing procedure can also be interpreted as an A^* search over a large hypergraph.

Each rule state `r` is characterized by the following fields:

- `rule` is the production rule currently being parsed.
- `start` is the (inclusive) sentence position marking the beginning of the production rule.
- `end` is the (exclusive) sentence position marking the end of the production rule.
- `i` is the current position in the sentence.
- `k` is the current position in the production rule. Dotted rule notation is a convenient way to represent the variables `rule` and `k`. For example, if the parser is currently examining the rule $A \to B_1 \dots B_n$ at rule position `k` (omitting semantic transformation operations), we may write this as $A \to B_1 \dots B_k \bullet B_{k+1} \dots B_n$ where the dot denotes the current position of the parser.
- `semantics` is a set of semantic statements.
- `syntax` is a partially completed syntax tree. As an example, if the parser is currently examining rule $A \to B_1 \dots B_n$ at position `k`, the tree will have a root node labeled A with `k` child subtrees each labeled B_1 through B_k, respectively.
- `log_probability` is the *inner* log probability of the rule up to its current position.

Every *complete* rule state `r` contains the above fields in addition to an `iterator` field which keeps intermediate state for the inference method described in section 3.1.1 (see description of the *iteration* operation below for details on how this is used).
Every nonterminal structure `n` contains the fields:

- `start`, `end`, `semantics`, `syntax`, and `log_probability` are identical to the respective fields in the rule states.
- `nonterminal` is the nonterminal currently being parsed.

The following are the available operations or deductions that the parser can perform while processing rule states:

expansion takes an incomplete rule state `r` as input. For notational convenience, let $k = \texttt{r.k}$ and `r.rule` be written as $A \to B_1 \dots B_n$. This operation examines the next right-hand symbol B_k. There are two possible cases:

If B_k is a nonterminal: For every production rule in the grammar $B_k \to \beta$ whose left-hand symbol is B_k, for every $j \in \{\texttt{r.i}, \dots, \texttt{r.end}\}$, create a new rule state $\texttt{r}^*$ *only if B_k was not previously expanded at this given start and end position*:

$\texttt{r}^*\texttt{.rule} = B_k \to \beta$, $\texttt{r}^*\texttt{.start} = \texttt{r.i}$, $\texttt{r}^*\texttt{.end} = j$,
$\texttt{r}^*\texttt{.i} = \texttt{r.i}$, $\texttt{r}^*\texttt{.k} = 0$, $\texttt{r}^*\texttt{.log_probability} = 0$.

The semantic statement field of the new state is set to be the set of all semantic statements for the expanded nonterminal: $\texttt{r}^*\texttt{.semantics} = \mathcal{X}_{B_k}$. The syntax tree of the new rule state $\texttt{r}^*\texttt{.syntax}$ is initialized as a single root node. The new rule state is added to the agenda (we address specifics on prioritization later). This operation is analogous to the "prediction" step in Earley parsing.

If B_k is a terminal: Read the terminal B_k in the sentence starting at position `r.i`, then create a new rule state r^* where:

r^*.rule = r.rule, r^*.start = r.start, r^*.end = r.end,
r^*.i = r.i + $|B_k|$, r^*.k = r.k + 1,
r^*.log_probability = r.log_probability,
r^*.semantics = r.semantics.

The new syntax tree is identical to the old syntax tree with an added child node corresponding to this terminal symbol B_k. The new rule state is then added to the agenda. This operation is analogous to the "scanning" step in Earley parsing.

completion takes as input an incomplete rule state `r`, and a nonterminal structure `n` where `n.nonterminal` matches the next right-hand nonterminal in `r.rule`, and where the starting position of the nonterminal structure `n.start` matches the current sentence position of the rule state `r.i`. For notational convenience, let `r.rule` be written as $A \rightarrow B_1{:}f_1 \ldots B_n{:}f_n$. The operation constructs a new rule state r^*:

r^*.rule = r.rule, r^*.start = r.start, r^*.end = r.end,
r^*.i = n.end, r^*.k = r.k + 1.

To compute the semantic statements of the new rule state, first invert the semantic statements of the nonterminal structure `n` with the semantic transformation operation f_k^{-1}, and then intersect the resulting set with the semantic statements of the incomplete rule state: r^*.semantics = r.semantics $\cap \{f_k^{-1}(x) : x \in$ n.semantics$\}$. The syntax tree of the new rule state r^*.syntax is the syntax tree of the old incomplete rule state `r.syntax` with the added subtree of the nonterminal structure `n.syntax`. The log probability of the new rule state is the sum of that of both input states: r^*.log_probability = r.log_probability + n.log_probability. The new rule state r^* is then added to the agenda. This operation is analogous to the "completion" step in Earley parsing.

iteration takes as input a complete rule state `r`. Having completed parsing the production rule r.rule = $A \rightarrow \beta$, we need to compute $\sum_{\theta_A^*} p(A \rightarrow \beta | x^{new}, \theta_A^*)$ as in equation 35. To do so, we determine HDP paths in order from highest to lowest posterior predictive probability using the HDP inference approach described in section 3.1.1. We store our current position in the list as `r.iterator`. This operation increments the iterator and adds the rule state back into the agenda (if the iterator has a successive element). Next, this operation creates a new nonterminal structure n^* where:

n^*.nonterminal = A, n^*.start = r.start,
n^*.end = r.end, n^*.syntax = r.syntax.

Recall that the paths in an HDP induce a partition of the set of semantic statements $\mathcal{X}_A$, and so the path returned by the iterator corresponds to a subset of semantic statements $X \subseteq \mathcal{X}_A$. The semantic statements of the nonterminal structure is computed as the intersection of this subset with semantic statements of the rule state: n^*.semantics = $X \cap$ r.semantics. The log probability of the new nonterminal structure n^*.log_probability is the sum of the log probability of the path returned by the iterator and `r.log_probability`. The new nonterminal structure is added to the chart.

The algorithm is started by executing the *expansion* operation on all production rules of the form $S \rightarrow \beta$ where S is the root nonterminal, starting at position 0 in the sentence, with `semantics` initialized as the set of all possible semantic statements $\mathcal{X}_S$.

To describe the prioritization of agenda items, recall that any complete syntax tree s can be decomposed into inner, left outer, and right outer portions: $s = s_L \cup s_R \cup s_I$. Observe that the probability of the full parse (equation 35) can be written as a product of four terms: (1) the semantic prior $p(x^{new})$, (2) the left outer probability $p(s_L|x^{new}, \boldsymbol{x}, \boldsymbol{s}, \boldsymbol{y})$, (3) the right outer probability $p(s_R|x^{new}, \boldsymbol{x}, \boldsymbol{s}, \boldsymbol{y})$, and (4) the inner probability $p(s_I|x^{new}, \boldsymbol{x}, \boldsymbol{s}, \boldsymbol{y})$.

Items in the agenda are sorted by an upper bound on the log probability of the entire parse. In order to compute this, we rely on an upper bound on the inner probability that only considers the syntactic structure:

$$\mathcal{I}_{\mathrm{A},i,j} \triangleq \max_{\mathrm{A}\to\mathrm{B}_1\ldots\mathrm{B}_n} \left(\max_{x'} \log p(\mathrm{A} \to \mathrm{B}_1 \ldots \mathrm{B}_n | x', \boldsymbol{x}, \boldsymbol{s}, \boldsymbol{y}) + \max_{m_2 \le \ldots \le m_n} \sum_{k=1}^{n} \mathcal{I}_{\mathrm{B}_k, m_k, m_{k+1}} \right), \tag{36}$$

where $m_1 = i$ and $m_{n+1} = j$. In the sum, if B_k is a terminal, then $\mathcal{I}_{\mathrm{B}_k,m_k,m_{k+1}} = 0$ if $m_{k+1} - m_k = |\mathrm{B}_k|$ is the correct length of the terminal; otherwise $\mathcal{I}_{\mathrm{B}_k,m_k,m_{k+1}} = -\infty$. The term $\max_{x'} \log p(\mathrm{A} \to \mathrm{B}_1 \ldots \mathrm{B}_n | x', \boldsymbol{x}, \boldsymbol{s}, \boldsymbol{y})$ can be computed exactly using algorithm 1, but a tight upper bound can be computed more quickly by terminating algorithm 1 early and using the priority value given by equation 25 (we find that for preterminals, even using the priority computed at the root provides a very good estimate). The value of $\mathcal{I}$ can be computed efficiently using existing syntactic (e.g., PCFG) parsers in time $\mathcal{O}(n^3)$.

We also compute an upper bound on the log probability of the outer portion of the syntax tree and the semantic prior. To be more precise, let $\mathcal{P}(\mathrm{A}, i, j)$ be the set of all parses (x, s_L, s_R, s_I) such that $s = s_L \cup s_R \cup s_I$ is the syntax and s_I is the inner syntax tree with root A that begins at sentence position i (inclusive) and ends at j (exclusive). Then, a bound on the outer probability is:

$$\mathcal{O}_{\mathrm{A},i,j} \triangleq \max_{(x,s_L,s_R,s_I)\in\mathcal{P}(\mathrm{A},i,j)} \log p(x) + \log p(s_L|x, \boldsymbol{x}, \boldsymbol{s}, \boldsymbol{y}) + \sum_{(\mathrm{A}',i',j')\in\mathcal{R}(s_R)} \mathcal{I}_{\mathrm{A}',i',j'}. \tag{37}$$

where $\mathcal{R}(s_R)$ is the set of root vertices of the trees contained in s_R, and $p(x)$ is the prior probability of the semantic statement $x \in \mathcal{X}_S$. Note that the third term is an upper bound on the right outer probability $p(s_R|x, \boldsymbol{s}, \boldsymbol{x}, \boldsymbol{y})$.

Using these bounds, we can compute an upper bound on the overall log probability of the parse for any state. For a rule state `r`, the search priority is given by:

$$\begin{aligned} &\texttt{r.log_probability} + \max_{m_{k+1}\le\ldots\le m_n} \sum_{l=k}^{n} \mathcal{I}_{\mathrm{B}_l,m_l,m_{l+1}} \\ &\quad + \min\left\{\log p(\texttt{r.semantics}), \mathcal{O}_{\mathrm{A},\texttt{r.start},\texttt{r.end}}\right\}, \end{aligned} \tag{38}$$

where $\mathrm{A} \to \mathrm{B}_1 \ldots \mathrm{B}_k$ is the currently-considered rule `r.rule`, $m_k =$ `r.i`, and $m_{n+1} =$ `r.end`. Note the first two terms constitute an upper bound on the inner probability of the nonterminal A, and the third term is an upper bound on the outer probability and semantic prior. The second term can be computed efficiently using dynamic programming. We further tighten this by adding a term that bounds the log probability of the rule $\log p(\mathrm{A} \to \mathrm{B}_1 \ldots \mathrm{B}_k | x', \boldsymbol{x}, \boldsymbol{s}, \boldsymbol{y})$. The items in the agenda are prioritized by this quantity. As long as the `log_probability` field remains exact, as it does in our approach, the overall search will yield exact outputs. The use of a syntactic parser

to compute a tigher bound on the outer probability in an A* parser is similar to the approach of Klein and Manning (2003b).

Naive computation of equation 37 is highly infeasible, as it would require enumerating all possible outer parses. However, we can rely on the fact that our search algorithm is *monotonic*: the highest score in the agenda never increases as the algorithm progresses. We prove monotonicity by induction on the number of iterations. For a given iteration i, by the inductive hypothesis, the parser has visited all reachable rule states with priority strictly larger than the priority of the current rule state. We will show that all new rule states added to the priority queue at iteration i must have priority at most equal to the priority of the current rule state. Consider each operation:

In the *expansion* operation, let $i = \texttt{r.i}$ and $k = \texttt{r.k}$. If the next right-hand side symbol B_k is a terminal, the new agenda item will have score at most that of the old agenda item, since $\texttt{r}^*\texttt{.log_probability} = \texttt{r.log_probability}$ and the sum of inner probability bounds in equation 38 cannot increase. If the B_k is a nonterminal, then we claim that any rule state created by this operation must have priority at most the priority of the old agenda item. Suppose to the contrary that there exists a $j \in \{i, \ldots, \texttt{r.end}\}$, a rule $\mathrm{B}_k \rightarrow \mathrm{C}_1 \ldots \mathrm{C}_u$, and $m'_2 \leq \ldots \leq m'_u$ such that:

$$\begin{aligned}
&\min\{\log p(\mathcal{X}_{\mathrm{B}_k}), \mathcal{O}_{\mathrm{B}_k,i,j}\} + \sum_{l=1}^{u} \mathcal{I}_{\mathrm{C}_l,m'_l,m'_{l+1}} \\
&\quad > \texttt{r.log_probability} + \max_{m_{k+1} \leq \ldots \leq m_n} \sum_{l=k}^{n} \mathcal{I}_{\mathrm{B}_l,m_l,m_{l+1}} \\
&\quad + \min\left\{\log p(\texttt{r.semantics}), \mathcal{O}_{\mathrm{A},\texttt{r.start},\texttt{r.end}}\right\},
\end{aligned}$$

where $m_k = m'_1 = i$, $m'_u = j$, and $m_{n+1} = \texttt{r.end}$. Note that the left-hand side is bounded above by $\mathcal{I}_{\mathrm{B}_k,i,j} + \mathcal{O}_{\mathrm{B}_k,i,j}$ which implies, by the definition of $\mathcal{O}_{\mathrm{B}_k,i,j}$, that there exists a parse $(x^*, s^*_L, s^*_R, s^*_I) \in \mathcal{P}(\mathrm{B}_k, i, j)$ such that:

$$\begin{aligned}
&\log p(x^*) + \log p(s^*_L | x^*, \boldsymbol{x}, \boldsymbol{s}, \boldsymbol{y}) + \mathcal{I}_{\mathrm{B}_k,i,j} + \sum_{(\mathrm{A}',i',j') \in \mathcal{R}(s^*_R)} \mathcal{I}_{\mathrm{A}',i',j'} \\
&\quad > \texttt{r.log_probability} + \max_{m_{k+2} \leq \ldots \leq m_n} \sum_{l=k}^{n} \mathcal{I}_{\mathrm{B}_l,m_l,m_{l+1}} \\
&\quad + \min\left\{\log p(\texttt{r.semantics}), \mathcal{O}_{\mathrm{A},\texttt{r.start},\texttt{r.end}}\right\},
\end{aligned}$$

where $m_{k+1} = j$. Let $\mathrm{C} \rightarrow \mathrm{D}_1 \ldots \mathrm{D}_v$ be the production rule in the syntax tree s^* containing s^*_I. In addition, let s^*_i be the sibling subtree of s_I rooted at D_i. This parse implies the existence of a rule state $\texttt{r}^*$ where $\texttt{r}^*\texttt{.rule} = (\mathrm{C} \rightarrow \mathrm{D}_1 \ldots \mathrm{D}_v)$, $\texttt{r}^*\texttt{.start}$ and $\texttt{r}^*\texttt{.end}$ are the start and end positions of the vertex corresponding to C, $\texttt{r}^*\texttt{.i} = \texttt{r.i}$, $\texttt{r}^*\texttt{.log_probability} = \sum_{i=1}^{\texttt{r}^*\texttt{.k}} \log p(s^*_i | x^*, \boldsymbol{x}, \boldsymbol{s}, \boldsymbol{y})$. The search priority of this rule state would be:

$$\begin{aligned}
&\sum_{i=1}^{\texttt{r}^*\texttt{.k}-1} \log p(s^*_i | x^*, \boldsymbol{x}, \boldsymbol{s}, \boldsymbol{y}) + \max_{m_{k+1} \leq \ldots \leq m_v} \sum_{l=\texttt{r}^*\texttt{.k}}^{v} \mathcal{I}_{\mathrm{D}_l,m_l,m_{l+1}} \\
&\qquad + \min\{\log p(\texttt{r}^*\texttt{.semantics}), \mathcal{O}_{\mathrm{C},\texttt{r}^*\texttt{.start},\texttt{r}^*\texttt{.end}}\}
\end{aligned}$$

We claim that this search priority must be strictly greater than that of the old agenda item `r`. By the definition of $\mathcal{O}$:

$$\mathcal{O}_{\mathrm{C},\mathtt{r}^*.\mathtt{start},\mathtt{r}^*.\mathtt{end}} \geq \log p(x^*) + \log p(s_L^* \setminus \{s_1^*, \ldots, s_{\mathtt{r}^*.\mathtt{k}-1}^*\} | x^*, \boldsymbol{x}, \boldsymbol{s}, \boldsymbol{y}) + \sum_{(\mathrm{A}',i',j') \in \mathcal{R}(s_R^* \setminus \{s_{\mathtt{r}^*.\mathtt{k}+1}^*, \ldots, s_v^*\})} \mathcal{I}_{\mathrm{A}',i',j'},$$

combined with the fact that $\log p(x^*) \leq \log p(X)$ for any $X \in \mathcal{X}_C$, observe that the search priority of $\mathtt{r}^*$ must be at least:

$$\log p(x^*) + \log p(s_L^* | x^*, \boldsymbol{x}, \boldsymbol{s}, \boldsymbol{y}) + \mathcal{I}_{\mathrm{B}_k,i,j} + \sum_{(\mathrm{A}',i',j') \in \mathcal{R}(s_R^*)} \mathcal{I}_{\mathrm{A}',i',j'},$$

which, in turn, is strictly greater than the priority of `r`. Thus, the priority of $\mathtt{r}^*$ is strictly larger than that of `r`, which would imply that the nonterminal B_k was previously expanded with start position i and end position j, which is a contradiction.

In the *completion* operation, the inner log probability of the new rule state $\mathtt{r}^*.\mathtt{log_probability}$ is at most the sum of the inner log probability of the old rule state $\mathtt{r}.\mathtt{log_probability}$ and the bound $\max_j \mathcal{I}_{\mathrm{B}_k,i,j}$. Thus, the priority of the new rule state is bounded by the priority of the old rule state.

In the *iteration* operation, monotonicity is guaranteed since the iterator structure returns items in order of non-increasing probability.

Therefore, the parser is monotonic. As a consequence, whenever the algorithm first expands a nonterminal B_k from a rule $\mathrm{A} \to \mathrm{B}_1 \ldots \mathrm{B}_n$, at start position i and end position j in the sentence, we have found the left outer parse that maximizes equation 37:

$$\mathcal{O}_{\mathrm{B}_{\mathtt{r}.\mathtt{k}},i,j} = \mathtt{r}.\mathtt{log_probability} + \max_{m_{k+2} \leq \cdots \leq m_n} \sum_{l=k+1}^{n} \mathcal{I}_{\mathrm{B}_l,m_l,m_{l+1}} + \min\left\{\log p(\mathtt{r}.\mathtt{semantics}), \mathcal{O}_{\mathrm{A},\mathtt{r}.\mathtt{start},\mathtt{r}.\mathtt{end}}\right\},$$

thereby computing $\mathcal{O}_{\mathrm{B}_{\mathtt{r}.\mathtt{k}},i,j}$ at no additional cost. Similarly, when the parser first constructs a nonterminal structure for the symbol A at start position i and end position j, monotonicity guarantees that no other nonterminal structure at (A, i, j) will have higher probability. We exploit this by updating the value of $\mathcal{I}_{\mathrm{A},i,j}$ as the algorithm progresses, incorporating more semantic information in the values of $\mathcal{I}$.

Figure 5
A step-by-step example of the parser running on the sentence "Chopin plays" using the grammar very similar to the one shown in figure 2. The top-left table lists the semantic statements sorted by their log probability of drawing the observation "Chopin" from the HDP associated with the nonterminal N. The top-center and top-right tables are defined similarly.

N → "Chopin"	log prob.
musician:chopin	-2
sport:swimming	-8
sport:tennis	-8
instrument:piano	-8
*	-8

The symbol * is a *wildcard*, referring to any entity in the ontology excluding those listed.
In this example, we use the grammar in figure 2.

V → "plays"	log prob.
athlete_plays_sport(, *)*	-2
musician_plays_inst(, *)*	-2
musician_plays_inst(, instrument:piano)*	-2
athlete_plays_sport(, sport:tennis)*	-2
athlete_plays_sport(, sport:swimming)*	-8
athlete_plays_sport(, instrument:piano)*	-8
⋮	⋮

VP → V	log prob.
athlete_plays_sport(, *)*	-4
musician_plays_inst(, *)*	-4
athlete_plays_sport(, sport:swimming)*	-4
musician_plays_inst(, instrument:piano)*	-5
athlete_plays_sport(, sport:tennis)*	-5
athlete_plays_sport(, instrument:piano)*	-8
⋮	⋮

iteration	*operation*	*new states created*
0	*expand* S at start position 0, end position 12	S → • N VP i: 0, end: 12 log_prob: 0
1	pop the S → • N VP state and *expand* N at start position 0 and end positions $0, \ldots, 12$	N → "Chopin" • i: 6 , end: 6 log_prob: 0
2	*iterate* the complete rule state N → "Chopin" •	N start: 0, end: 6 log_prob: -2 *musician:chopin*
3	pop the N nonterminal state and *complete* any waiting rule states	S → N • VP i: 7, end: 12 log_prob: -2 * (*musician:chopin, **)
4	pop the S → N • VP state and *expand* VP at start position 7 and end position 12	VP → • V i: 7, end: 12 log_prob: 0 VP → • V N i: 7, end: 12 log_prob: 0 VP → • V i: 7, end: 11 log_prob: 0 . . .
5	pop the VP → • V state and *expand* V	V → "plays" • i: 12, end: 12 log_prob: 0
6	*iterate* the complete rule state V → "plays" •	V start: 7, end: 12 log_prob: -2 *athlete_plays_sport(*, *)* V start: 7, end: 12 log_prob: -2 *musician_plays_inst(*, *)* . . .
7	pop a V nonterminal state and *complete* any waiting rule states	VP → V • i: 12, end: 12 log_prob: -2 *musician_plays_inst(*, *)* VP → V • N i: 12, end: 12 log_prob: -2 *musician_plays_inst(*, *)* . . .
8	*iterate* the complete rule state VP → V •	VP start: 7, end: 12 log_prob: -6 *musician_plays_inst(*, *)*
9	pop the VP nonterminal state and *complete* any waiting rule states	S → N VP • i: 12, end: 12 log_prob: -8 *musician_plays_inst(* *musician:chopin, *)*

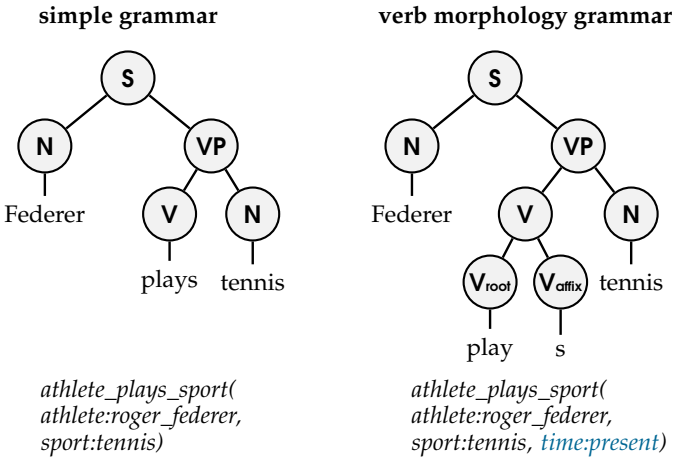

Figure 6
An example of a (simpified) labeled data instance in our experiments. For brevity, we omit semantic transformation operations, syntax elements such as word boundaries, irregular verb forms, etc.

6. Results

The experiments in this section evaluate our parser's ability to parse semantic statements from short sentences, consisting of a subject noun, a simple verb phrase, and an object noun. We also evaluate the ability to incorporate background knowledge during parsing, through the semantic prior. To do so, we used the ontology and knowledge base of the Never-Ending Language Learning system (NELL) (Mitchell et al. 2015). We use a snapshot of NELL at iteration 905 containing 1,786,741 concepts, 623 relation predicates, and 2,212,187 beliefs (of which there are 131,365 relation instances). The relations in NELL are *typed*, where the domain and range of each relation is a category in the ontology. We compare our parser to a state-of-the-art CCG parser (Krishnamurthy and Mitchell 2014) trained and tested on the same data.

6.1 Relation extraction

We first evaluate our parser on a relation extraction task on a dataset of subject-verb-object (SVO) sentences. We created this dataset by filtering and labeling sentences from a corpus of SVO triples (Talukdar, Wijaya, and Mitchell 2012) extracted from dependency parses of the ClueWeb09 dataset (Callan et al. 2009). NELL provides a `can_refer_to` relation, mapping noun phrases to concepts in the NELL ontology. We created our own mapping between verbs (or simple verb phrases) and 223 relations in the NELL ontology. Using these two mappings, we can identify whether an SVO triple can refer to a belief in the NELL knowledge base. We only accepted sentences that referred to high-confidence beliefs in NELL (for which NELL gives a confidence score of at least 0.999). The accepted sentences were labeled with the referred beliefs. For this experiment, we restrict all verbs to the present tense. This yielded a final dataset of 2,546 SVO three-word sentences, along with their corresponding semantic statement from the NELL KB, spanning over 74 relations and 1,913 concepts. We randomly split the data into a training set of 2,025 sentences and a test set of 521 sentences. In the task, each parser makes predictions on every test sentence, which we mark as correct if the

output semantic statement exactly matches the label. The main difficulty in this task is to learn the mapping between relations and the sentence text. For example, the dataset contains verbs such as 'makes' which can refer to at least five NELL relations, including *companyeconomicsector*, *directordirectedmovie* and *musicartistgenre*. The semantic types of the subject and object concepts are very informative in resolving ambiguity, and prior knowledge in the form of a belief system can further aid parsing. The precision-recall curves in figure 7 were generated by sorting the outputs of our parser by posterior probability, which was computed using the top $k = 10000$ output parses for each test sentence (see section 6.4 for experiments with varying k).

We call a semantic statement "type-correct" if the subject and object concepts agree with the domain and range of the instantiated relation, under the NELL ontology. We experimented with three prior settings for our parser: (1) uniform prior, (2) a prior where all type-correct semantic statements have a prior probability that is larger by 4 units (in terms of log probability) than type-incorrect statements, and (3) a prior where all semantic statements that correspond to true beliefs in the KB have a prior probability that is 8 larger than type-incorrect statements and all type-correct correct statements have probability 4 larger than type-incorrect statements.

In the simple relation extraction task, we find that CCG performs comparably to our parser under a uniform and type-correct prior. In fact, the parsers make the almost identical predictions on the test sentences. The differences in the precision-recall curves arise due to the differences in the scoring of predictions. The primary source of incorrect predictions is when a noun in the test set refers to a concept in the ontology but does not refer to the same concept in the training set. For example, in the sentence "Wilson plays guitar," both parsers predict that "Wilson" refers to the politician Greg Wilson. The similarity in the performance of our parser with the uniform prior and the type-correct prior suggests that the parser learns "type-correctness" from the training data. This is due to the fact that, in our grammar, the distribution of the verb depends jointly on the types of both arguments. With the KB prior, our parser outperforms CCG, demonstrating that our parser effectively incorporates background knowledge via the semantic prior to improve precision and recall.

6.2 Modeling word morphology

In the second experiment, we demonstrate our parser's ability to extract semantic information from the morphology of individual verbs, by operating over characters instead of preprocessed tokens. We generated a new labeled SVO dataset using a process similar to that in the first experiment. In this experiment, we did not restrict the verbs to the present tense. This dataset contains 3,197 sentences spanning 56 relations and 2,166 concepts. The data was randomly split into a set of 2,538 training sentences and 659 test sentences. We added a simple temporal model to the semantic formalism: all sentences in any past tense refer to semantic statements that were true in the past; all sentences in any present tense refer to presently true statements; and sentences in any future tense refer to statements that will be true in the future. Thus the task becomes one of *temporally-scoped relation extraction*. A simple verb morphology model was incorporated into the grammar. Each verb is modeled as a concatenated root and affix. In the grammar, the random selection of a production rule captures the selection of the verb tense. The affix is selected deterministically according to the desired tense and the grammatical person of the subject. The posterior probability of each parse was estimated using the top $k = 10000$ parses for each test example. Results are shown in figure 8.

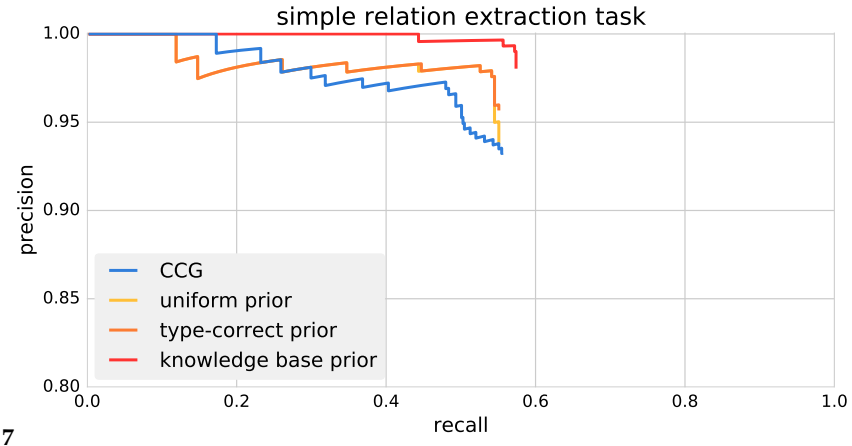

Figure 7
Precision-recall curves for the standard relation extraction task. We compare our approach with three different settings of the prior (yellow, orange, red) and CCG (blue). The uniform prior places equal probability mass on all semantic statements. The "type-correct" prior places higher mass (+4 in log probability) on semantic statements with subject and object types that agree with the domain and range of the relation predicate. The "knowledge base prior" is similar to the type-correct prior, except that it additionally places higher probability mass on semantic statements that correspond to true beliefs in the NELL knowledge base (an additional +4 in log probability).

Table 1
A sample of two randomly selected parses from the simple relation extraction task, using a uniform prior and $k = 10^7$. For each of the two sample sentences, the top few parse outputs are displayed, along with their log probabilities. Recall that our parser operates over *sets* of semantic statements, so some of the outputs contain wildcards. It is evident from the output on the right that the phrase "Kidneys" did not appear in the training set, and so the highest-ranked parse is ambiguous.

"Mickey Rourke stars in Wrestler"	log prob.
actor_starred_in_movie(*actor:mickey_rourke,* *movie:wrestler)*	-10.50
actor_starred_in_movie(any *actor* except *actor:miley_cyrus,* *actor:mike_myers, etc.,* *movie:wrestler)*	-15.11
actor_starred_in_movie(*actor:mickey_rourke,* any *movie* except *movie:austin_powers,* *movie:cold_mountain,* *etc.)*	-15.11
actor_starred_in_movie(*actor:anne_hathaway,* *movie:wrestler)*	-15.11
⋮	⋮

"Kidneys contain blood vessels"	log prob.
bodypart_contains_bodypart(any *bodypart* except *artery:hand,* *braintissue:brains, etc.,* *artery:blood_vessels)*	-15.11
bodypart_contains_bodypart(*bodypart:blood,* *artery:blood_vessels)*	-15.11
bodypart_contains_bodypart(*bodypart:legs,* *artery:blood_vessels)*	-15.11
bodypart_contains_bodypart(*braintissue:parts,* *artery:blood_vessels)*	-15.11
bodypart_contains_bodypart(*bodypart:nerves001,* *artery:blood_vessels)*	-15.11
⋮	⋮

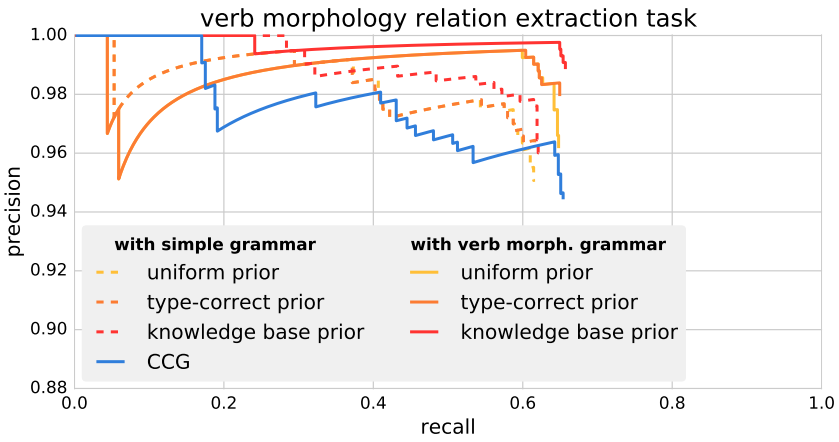

Figure 8
Precision-recall curves for the temporally-scoped relation extraction task. We compare our approach with three different settings of the prior (yellow, orange, red) and CCG (blue). The solid lines denote our parser's performance when using a grammar that models the morphology of verbs, whereas the dashed lines are produced when our parser is trained with a grammar that does not model verb morphology. The uniform prior places equal probability mass on all semantic statements. The "type-correct" prior places higher mass (+4 in log probability) on semantic statements with subject and object types that agree with the domain and range of the relation predicate. The "knowledge base prior" is similar to the type-correct prior, except that it additionally places higher probability mass on semantic statements that correspond to true beliefs in the NELL knowledge base (additional +4 in log probability).

In this temporally-scoped relation extraction task, our parser demonstrates better generalization over verb forms. Our parser performs better when trained on a grammar that models verb morphology (solid lines) than when trained on a simpler grammar that does not consider the morphology of verbs (dashed lines). The parser is able to accomplish this due to its ability to model the semantics in the morphology of individual words. As in the first experiment, the performance of our parser improves when we use the knowledge base prior, supporting the observation that our parser can effectively leverage background knowledge to improve its accuracy. There is, again, no difference in performance when using the uniform prior vs. the type-correct prior.

6.3 Out-of-vocabulary behavior

Recall that our parser operates over sets of semantic statements, as opposed to individual statements. Thus, it is possible that when parsing completes, the output is a non-singleton set of semantic statements that all share the highest probability parse. In our first two experiments, we counted these outputs as a "non-parse" to more fairly compare with CCG. However, we can evaluate the quality of these *ambiguous* outputs. In figure 9, we again perform the simple relation extraction task with a modification: we measure the correctness of our parser's output by whether the ground truth semantic statement is contained within the set of semantic statements that share the highest probability parse. CCG does not produce output on sentences that contain tokens which did not appear in the training data. In a sense, this evaluation measures out-of-vocabulary performance. Although precision is not as high as the in-vocabulary test, recall is much

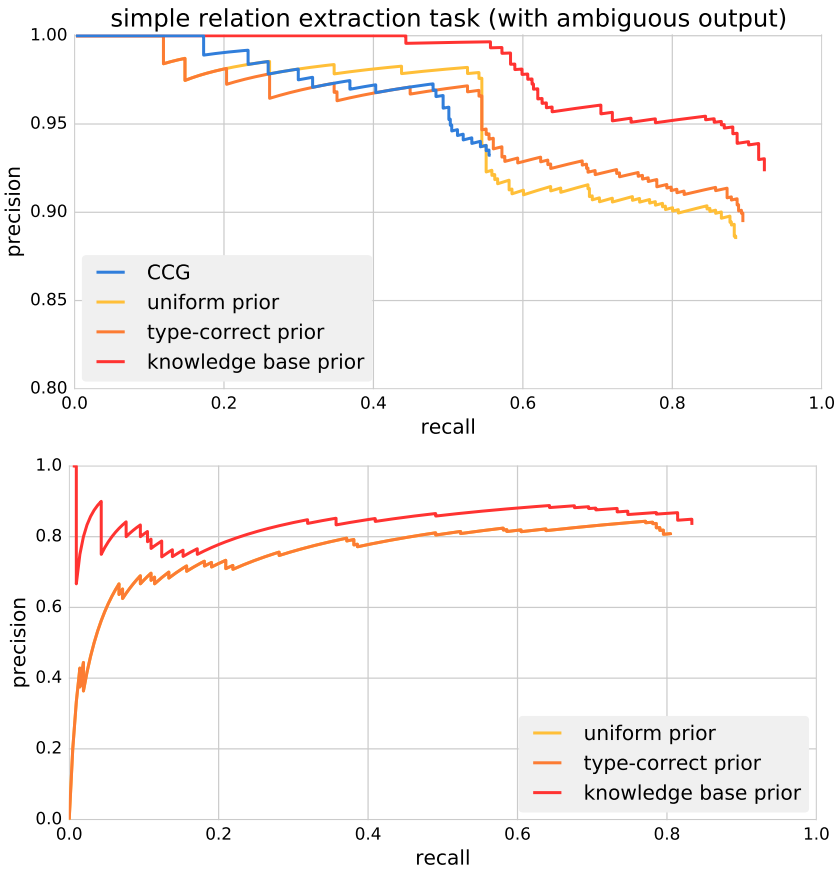

Figure 9
(**top**) Precision and recall for the simple relation extraction task, including unambiguous outputs, and (**bottom**) on the subset of sentences for which CCG did not provide output. We compare our approach with three different settings of the prior (yellow, orange, red) and CCG (blue). The uniform prior places equal probability mass on all semantic statements. The "type-correct" prior places higher mass (+4 in log probability) on semantic statements with subject and object types that agree with the domain and range of the relation predicate. The "knowledge base prior" is similar to the type-correct prior, except that it additionally places higher probability mass on semantic statements that correspond to true beliefs in the NELL knowledge base (additional +4 in log probability).

improved. The knowledge base prior again results in improved performance over the less informative priors.

6.4 Effect of changing the parameter k

In our parser, the search algorithm does not terminate until it has found the k-best semantic parses. This is useful for evaluating the confidence of the parser in its output, and for estimating the posterior probability of each parse. We examine the behavior of the parser as a function of k in figure 10. The timing results demonstrate that the parser can scale to large knowledge bases while maintaining efficiency during parsing.

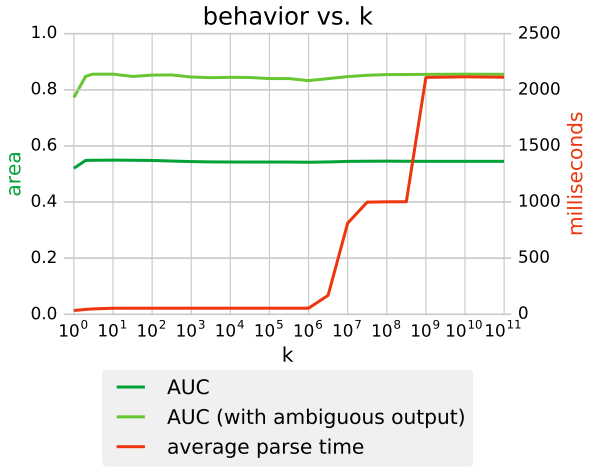

Figure 10
Area under the precision-recall curve and average parse time versus k, for the simple relation extraction task and a uniform prior. The dark green curve measures the area under the precision-recall curve without considering ambiguous outputs (as in figure 7), whereas the light green curve measures the area under the precision-recall curve taking into account the ambiguous outputs (as in the top plot in figure 9)

Recall that, in order to produce the precision-recall curves for our parser's output, we sort the outputs by their confidence scores (estimates of the posterior probability). In the figure, the area under the precision-recall curve (AUC) converges quickly at fairly small values of k, indicating that the relative ordering of the parse outputs converges quickly. We found this behavior to be consistent for other choices of prior distributions and for the more complex temporally-scoped relation extraction task. For values of k smaller than $\sim 10^6$, the parser provides outputs very quickly. This is due to the fact that, as the parser performs its search, it can quickly find a set of semantic statements of size roughly 10^6 (e.g., when parsing "Federer plays golf", it will quickly find an output that looks like *athlete_plays_sport(athlete:roger_federer,*)* where * is a wildcard that denotes any concept in the ontology). The parser will require additional time to search beyond this initial ambiguous output. Note that this threshold value of k is identical to the number of concepts in the ontology: 1,786,741. The second threshold is likely related to the product of the number of concepts and the number of relations in the ontology: $623 \cdot 1786741 \approx 1.11 \cdot 10^9$.

6.5 Out-of-knowledge base relation extraction

In all earlier experiments, we used a dataset that only contained sentences that refer to beliefs in the NELL knowledge base. In order to inspect the performance of the parser on sentences that do not refer to NELL beliefs, we create a new dataset. We again start with the SVO sentence corpus and modify the filtering process: we only accept sentences that (1) contain noun phrases that also exist in the first dataset (to ensure the parsers at least have a chance to unambiguously understand the sentences), (2) contain verb phrases that exist in the hand-constructed verb-relation map we used to create the first dataset, (3) cannot refer to any NELL belief, according to the `can_refer_to` instances and the verb-relation map. More precisely, for every sentence, the `can_refer_to` relation

Table 2
Precision and recall of the parsers evaluated on the out-of-knowledge base dataset. The uniform prior places equal probability mass on all semantic statements. The "type-correct" prior places higher mass (+4 in log probability) on semantic statements with subject and object types that agree with the domain and range of the relation predicate. The "knowledge base prior" is similar to the type-correct prior, except that it additionally places higher probability mass on semantic statements that correspond to true beliefs in the NELL knowledge base (additional +4 in log probability).

simple relation extraction	**precision**	**recall**	**F1**
CCG	0.65	0.44	0.52
our parser (no prior)	0.78	**0.58**	**0.66**
our parser (type-correct prior)	0.84	0.46	0.60
our parser (KB prior)	**0.90**	0.45	0.60
temporally-scoped relation extr.	**precision**	**recall**	**F1**
CCG	0.78	0.63	0.69
our parser (no prior)	0.81	**0.68**	**0.73**
our parser (type-correct prior)	**0.84**	0.59	0.69
our parser (KB prior)	**0.84**	0.54	0.65

maps each noun phrase to a set of possible referent concepts; the relation-verb map provides a set of possible referent relations; and so their Cartesian product provides a set of possible referent semantic statements. We discard those sentences where the set of referent semantic statements contains a NELL belief. We sorted the resulting sentence list by frequency and labeled them by hand. Of the 1365 most frequent sentences in the filtered set, we labeled 100, since some sentences referred to concepts outside of the ontology or their verbs referred to unrecognized relations (i.e., not in the verb-relation map). This dataset is referred to as the *out-of-knowledge base* dataset, since the sentences refer to beliefs outside the knowledge base. We selected 20 sentences from this dataset as training sentences. We trained all parsers on these sentences in addition to the entirety of the first dataset. We tested the parsers on the remaining 80 sentences.

Table 2 displays the performance results of our parser and CCG on this out-of-knowledge base dataset, in both the simple relation extraction task as well as the temporally-scoped task. Our parser is trained with the simple grammar in the simple relation extraction task, and with the grammar that models verb morphology in the temporally-scoped relation extraction task. As expected, the more informative priors do not uniformly improve parsing performance in this evaluation. Interestingly, the parser behaves more conservatively when incorporating the stronger priors, outputting a smaller set of the most confident responses, which results in higher precision and reduced recall. Our parser is indeed capable of extracting the correct semantic statements from sentences that refer to beliefs outside the knowledge base, and the use of the informative priors does not obviously hurt performance.

7. Discussion

In this article, we presented a generative model of sentences for semantic parsing, extending the CFG formalism to couple semantics with syntax. In the generative process, a semantic statement is first generated, for example, from a knowledge base. Next,

the tree is constructed top-down using a recursive procedure: production rules are selected randomly, possibly depending on features of the semantic statement. Semantic transformation operations specify how to decompose the semantic statement in order to continue recursion. We presented a particular construction where production rules are selected using an HDP. We applied MCMC to perform inference in this model, and constructed a chart-driven agenda parser. Our application of the HDP is distinct from previous uses, since in our construction, the path indicator for each observation is not assumed to be fixed. We evaluate our parser on a dataset of SVO sentences, labeled with semantic statements from NELL. The results demonstrate that our parser can incorporate prior knowledge from a knowledge base via the semantic prior. With an informative prior, our parser outperforms a state-of-the-art CCG parser. In addition, we demonstrate that our model can be used to jointly model the morphology of individual verbs, leading to improved generalization over verbs in a temporally-scoped relation extraction task. The results indicate that our framework can scale to knowledge bases, such as NELL, with millions of beliefs, and can be extended to more complex grammars and richer semantic formalisms without sacrificing exact inference and the principled nature of the model.

An interesting parallel can be drawn between our inference problem and the problem of finding shortest paths in hypergraphs. Similar parallels have been drawn in other parsers (Klein and Manning 2001, 2003a; Pauls and Klein 2009; Pauls, Klein, and Quirk 2010). Since our approach is top-down, the specification of our hypergraph is more involved. Imagine a hypergraph containing a vertex for every semantic statement $x \in \mathcal{X}$, a vertex for every intermediate rule state $\mathrm{A} \to \mathrm{B}_1 \dots \mathrm{B}_k \bullet \mathrm{B}_{k+1} \dots \mathrm{B}_n$, and two vertices for every nonterminal (one indicating that parsing is incomplete and one for completed parses). Add a hyperedge to this graph for every allowable operation by the parser. A hyperedge is a generalization of an edge where both its "head" and "tail" can be sets of vertices. Then, the problem of parsing can be equivalently stated as finding the shortest "path" from two sets of vertices: the source set of vertices are those representing the incomplete S nonterminal and all elements of $\mathcal{X}_S$, and the destination vertex is the complete S nonterminal. See Gallo, Longo, and Pallottino (1993), Klein and Manning (2001) for definitions and further details. Our algorithm can then be understood as an application of A^* search for the k-best paths in this hypergraph. The monotonicity property of our algorithm is a consequence of Dijkstra's theorem generalized to hypergraphs (Gallo, Longo, and Pallottino 1993). This also suggests that our parser can be improved by utilizing a tighter heuristic.

In the parser, the prior contribution is rather loosely incorporated into the objective (equation 38). This is due to the fact that we assumed nothing about the structure of the semantic prior. However, the algorithm could potentially be made more efficient if we could factorize the prior, for example, over nonterminals or rules in the grammar. This could provide an additive term to equation 38.

We presented our inference approach as a combination of two search algorithms: (1) the HDP inference component, and (2) the joint syntactic-semantic parser. However, it is possible to merge these two searches into a single search (the priority queues of each algorithm would be unified into a single global priority queue), potentially improving the overall efficiency.

In this article, we showed how to utilize HDPs to add dependence between semantic features and the probabilistic selection of production rules in the generative process. It would be interesting to explore the application of other dependent Dirichlet processes and random probability measures, possibly as a means to induct a grammar in our framework.

Another highly promising avenue of research is to explore more complex prior structures. For instance, a generative model of a knowledge base could be composed with the framework we present here. This would result in a parser that would learn new beliefs as it reads text. Another direction is to model the generation of a sequence of sentences (such as in a paragraph), where complex relationships between concepts bridge across multiple sentences. Such an approach would likely contain a model of *context* shared across sentences. An example of such a generative process would be to "generate" a representation of the document context using a background knowledge base. Then, semantic statements for each sentence in the document can be generated from this intermediate document-level representation in addition to other sources of document-level information. Finally, our grammar would generate sentences for each semantic statement. The problem of co-reference resolution also becomes more apparent in these settings with more complex sentences. In both the single-sentence and multiple-sentence settings, co-reference resolution can be integrated into parsing. Incorporating these extensions into a richer unified parsing framework would be promising.

Acknowledgments

We thank Emmanouil A. Platanios and Jayant Krishnamurthy for the insightful discussion and helpful comments on the draft. We also thank anonymous reviewers for their feedback on the shortened version of this manuscript. This work was supported in part by NSF grant IIS1250956, and in part by DARPA DEFT contract FA87501320005.

References

Adams, Beverly Colwell, Laura C. Bell, and Charles A. Perfetti. 1995. A trading relationship between reading skill and domain knowledge in children's text comprehension. *Discourse Processes*, 20(3):307–323. 2

Aho, Albert V. and Jeffery D. Ullman. 1972. *The Theory of Parsing, Translation, and Compiling*, volume 1. Prentice-Hall, Englewood Cliffs, NJ. 4

Aldous, David J. 1985. Exchangeability and related topics. In *École d'été de probabilités de Saint-Flour, XIII—1983*, volume 1117 of *Lecture Notes in Math.* Springer, Berlin, pages 1–198. 5

Anderson, Richard C. and P. David Pearson. 1984. *A schema-theoretic view of basic processes in reading comprehension*. University of Illinois at Urbana-Champaign Champaign, Ill. : Bolt, Beranek, and Newman Inc. 2

Baker, Collin F., Charles J. Fillmore, and John B. Lowe. 1998. The berkeley FrameNet project. In *Proceedings of the 17th International Conference on Computational Linguistics*, volume 1, pages 86–90, Montreal, Quebec, Canada. Association for Computational Linguistics. 3

Banarescu, Laura, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, and Nathan Schneider. 2013. Abstract meaning representation for sembanking. 3

Berger, Adam, Stephen Della Pietra, and Vincent Della Pietra. 1996. A maximum entropy approach to natural language processing. *Computational Linguistics*, 22:39–71. 12

Bloom, Paul. 2000. *How Children Learn the Meanings of Words*. MIT Press, Cambridge, MA. 2

Bordes, Antoine, Xavier Glorot, Jason Weston, and Yoshua Bengio. 2012. Joint learning of words and meaning representations for open-text semantic parsing. In Neil D. Lawrence and Mark Girolami, editors, *AISTATS*, volume 22 of *JMLR Proceedings*, pages 127–135. JMLR.org. 1

Callan, Jamie, Mark Hoy, Changkuk Yoo, and Le Zhao. 2009. The clueweb09 dataset. 25

Chomsky, N., editor. 1957. *Syntactic structures*. Mouton & Co., The Hague. 4

Chomsky, Noam. 1956. Three models for the description of language. *IRE Transactions on Information Theory*, 2:113–124. 4, 11

Church, A. 1932. A set of postulates for the foundation of logic part i. *Annals of Mathematics*, 33(2):346–366. 3

Cohen, Shay B., David M. Blei, and Noah A. Smith. 2010. Variational inference for adaptor grammars. In *HLT-NAACL*, pages 564–572. The Association for Computational Linguistics. 4

Durrett, Greg and Dan Klein. 2014. A joint model for entity analysis: Coreference, typing, and linking. *TACL*, 2:477–490. 5
Earley, Jay. 1970. An efficient context-free parsing algorithm. *Communications of the ACM*, 6(8):451–455. Reprinted in Grosz et al. (1986). 19
Ferguson, Thomas S. 1973. A Bayesian analysis of some nonparametric problems. *The Annals of Statistics*, 1(2):209–230. 5
Fincher-Kiefer, Rebecca. 1992. The role of prior knowledge in inferential processing. *Journal of Research in Reading*, 15:12–27. 2
Finkel, Jenny Rose and Christopher D. Manning. 2009. Joint parsing and named entity recognition. In *HLT-NAACL*, pages 326–334. The Association for Computational Linguistics. 1
Finkel, Jenny Rose and Christopher D. Manning. 2010. Hierarchical joint learning: Improving joint parsing and named entity recognition with non-jointly labeled data. In Jan Hajic, Sandra Carberry, and Stephen Clark, editors, *ACL*, pages 720–728. The Association for Computer Linguistics. 1
Gallo, Giorgio, Giustino Longo, and Stefano Pallottino. 1993. Directed hypergraphs and applications. *Discrete Applied Mathematics*, 42(2):177–201. 4, 32
Gelfand, Alan E. and Adrian F. M. Smith. 1990. Sampling-based approaches to calculating marginal densities. *Journal of the American Statistical Association*, 85(410):398–409. 6
Hockenmaier, Julia. 2001. Statistical parsing for ccg with simple generative models. In *ACL (Companion Volume)*, pages 7–12. CNRS, Toulose, France. 12
Hockenmaier, Julia and Mark Steedman. 2002. Generative models for statistical parsing with combinatory categorial grammar. In *ACL*, pages 335–342. ACL. 12
Jäger, Gerhard. 2004. *Anaphora and type logical grammar*. UniversitÃd'tsbibliothek Johann Christian Senckenberg. 12
Johnson, Mark and Katherine Demuth. 2010. Unsupervised phonemic chinese word segmentation using adaptor grammars. In Chu-Ren Huang and Dan Jurafsky, editors, *COLING*, pages 528–536. Tsinghua University Press. 4
Johnson, Mark, Thomas L. Griffiths, and Sharon Goldwater. 2007. Adaptor grammars: A framework for specifying compositional nonparametric Bayesian models. In *Advances in Neural Information Processing Systems (NIPS)*. 4
Kazama, Jun'ichi and Kentaro Torisawa. 2007. Exploiting wikipedia as external knowledge for named entity recognition. In *Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 698–707. 1
Kim, J. and D. Moldovan. 1995. Acquisition of linguistic patterns for knowledge-based information extraction. *IEEE Transactiops on Knowledge and Data Engineering*, 7(5):713–724. 5
Klein, Dan and Christopher D. Manning. 2001. Parsing and hypergraphs. In *IWPT*. Tsinghua University Press. 4, 32
Klein, Dan and Christopher D. Manning. 2003a. A* parsing: Fast exact viterbi parse selection. In *HLT-NAACL*. 4, 32
Klein, Dan and Christopher D. Manning. 2003b. Fast exact inference with a factored model for natural language parsing. In *In Advances in Neural Information Processing Systems 15 (NIPS)*. 22
Krishnamurthy, Jayant and Tom M. Mitchell. 2014. Joint syntactic and semantic parsing with combinatory categorial grammar. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1188–1198, Baltimore, Maryland, June. Association for Computational Linguistics. 5, 25
Li, Junhui, Muhua Zhu, Wei Lu, and Guodong Zhou. 2015. Improving semantic parsing with enriched synchronous context-free grammar. In Lluís Màrquez, Chris Callison-Burch, Jian Su, Daniele Pighin, and Yuval Marton, editors, *EMNLP*, pages 1455–1465. The Association for Computational Linguistics. 4
Liang, Percy, Michael I. Jordan, and Dan Klein. 2013. Learning dependency-based compositional semantics. *Computational Linguistics*, 39(2):389–446. 3
Merlo, Paola and Gabriele Musillo. 2008. Semantic parsing for high-precision semantic role labelling. In Alexander Clark and Kristina Toutanova, editors, *CoNLL*, pages 1–8. ACL. 1
Mitchell, T., W. Cohen, E. Hruscha, P. Talukdar, J. Betteridge, A. Carlson, B. Dalvi, M. Gardner, B. Kisiel, J. Krishnamurthy, N. Lao, K. Mazaitis, T. Mohammad, N. Nakashole, E. Platanios, A. Ritter, M. Samadi, B. Settles, R. Wang, D. Wijaya, A. Gupta, X. Chen, A. Saparov, M. Greaves, and J. Welling. 2015. Never-ending learning. In *AAAI*. : Never-Ending Learning in AAAI-2015. 25

Nakashole, Ndapandula and Tom M. Mitchell. 2015. A knowledge-intensive model for prepositional phrase attachment. In *ACL (1)*, pages 365–375. The Association for Computer Linguistics. 5

Ng, Vincent. 2007. Shallow semantics for coreference resolution. In *IJCAI*, pages 1689–1694. 1

Pauls, Adam and Dan Klein. 2009. K-best a* parsing. In Keh-Yih Su, Jian Su, and Janyce Wiebe, editors, *ACL/IJCNLP*, pages 958–966. The Association for Computer Linguistics. 4, 32

Pauls, Adam, Dan Klein, and Chris Quirk. 2010. Top-down k-best a* parsing. In *Proceedings of the ACL 2010 Conference Short Papers*, ACLShort '10, pages 200–204, Stroudsburg, PA, USA. Association for Computational Linguistics. 4, 32

Ponzetto, S. P. and M. Strube. 2006. Exploiting semantic role labeling, WordNet and Wikipedia for coreference resolution. *Proceedings of the main conference on Human Language Technology Conference of the North American Chapter of the Association of Computational Linguistics*, pages 192–199. 1

Prokofyev, Roman, Alberto Tonon, Michael Luggen, Loic Vouilloz, Djellel Eddine Difallah, and Philippe CudrÃł-Mauroux. 2015. Sanaphor: Ontology-based coreference resolution. In Marcelo Arenas, ÃŞscar Corcho, Elena Simperl, Markus Strohmaier, Mathieu d'Aquin, Kavitha Srinivas, Paul T. Groth, Michel Dumontier, Jeff Heflin, Krishnaprasad Thirunarayan, and Steffen Staab, editors, *International Semantic Web Conference (1)*, volume 9366 of *Lecture Notes in Computer Science*, pages 458–473. Springer. 5

Ratinov, Lev-Arie and Dan Roth. 2012. Learning-based multi-sieve co-reference resolution with knowledge. In Jun'ichi Tsujii, James Henderson, and Marius Pasca, editors, *EMNLP-CoNLL*, pages 1234–1244. ACL. 5

Ratnaparkhi, Adwait. 1998. Maximum entropy models for natural language ambiguity resolution. *PhD thesis*. 12

Robert, C. P. and G. Casella. 2010. *Monte Carlo Statistical Methods*. Springer, New York, NY. 6

Salloum, Wael. 2009. A question answering system based on conceptual graph formalism. volume 3. IEEE CS Press. 5

Salton, Gerald. 1971. *The SMART Retrieval System: Experiments in Automatic Document Processing*. Prentice Hall. 3

Steedman, Mark. 1996. *Surface structure and interpretation*. Linguistic inquiry monographs, 30. MIT Press. 12

Talukdar, Partha Pratim, Derry Tanti Wijaya, and Tom M. Mitchell. 2012. Acquiring temporal constraints between relations. In Xue wen Chen, Guy Lebanon, Haixun Wang, and Mohammed J. Zaki, editors, *CIKM*, pages 992–1001. ACM. 25

Tanaka, Takaaki, Francis Bond, Timothy Baldwin, Sanae Fujita, and Chikara Hashimoto. 2007. Word sense disambiguation incorporating lexical and structural semantic information. In *EMNLP-CoNLL*, pages 477–485. ACL. 1

Teh, Yee Whye, Michael I. Jordan, Matthew J. Beal, and David M. Blei. 2006. Hierarchical dirichlet processes. *Journal of the American Statistical Association*, 101(476):1566–1581. 4, 6, 7, 10

Turney, Peter D. and Patrick Pantel. 2010. From frequency to meaning: Vector space models of semantics. *CoRR*, abs/1003.1141. 3

Wong, Yuk Wah and Raymond J. Mooney. 2006. Learning for semantic parsing with statistical machine translation. In Robert C. Moore, Jeff A. Bilmes, Jennifer Chu-Carroll, and Mark Sanderson, editors, *HLT-NAACL*. The Association for Computational Linguistics. 4

Wong, Yuk Wah and Raymond J. Mooney. 2007. Learning synchronous grammars for semantic parsing with lambda calculus. In John A. Carroll, Antal van den Bosch, and Annie Zaenen, editors, *ACL*. The Association for Computational Linguistics. 4